53-1002335-01
30 May 2012

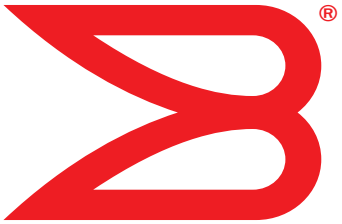

ServerIron Traffic Works

Firewall and Load Balancing Guide

Supporting ServerIron TrafficWorks version 10.2.02

BROCADE

Brocade Communications Systems, Incorporated

Corporate and Latin American Headquarters
Brocade Communications Systems, Inc.
130 Holger Way
San Jose, CA 95134
Tel: 1-408-333-8000
Fax: 1-408-333-8101
E-mail: info@brocade.com

Asia-Pacific Headquarters
Brocade Communications Systems China HK, Ltd.
No. 1 Guanghua Road
Chao Yang District
Units 2718 and 2818
Beijing 100020, China
Tel: +8610 6588 8888
Fax: +8610 6588 9999
E-mail: china-info@brocade.com

European Headquarters
Brocade Communications Switzerland Sàrl
Centre Swissair
Tour B - 4ème étage
29, Route de l'Aéroport
Case Postale 105
CH-1215 Genève 15
Switzerland
Tel: +41 22 799 5640
Fax: +41 22 799 5641
E-mail: emea-info@brocade.com

Asia-Pacific Headquarters
Brocade Communications Systems Co., Ltd. (Shenzhen WFOE)
Citic Plaza
No. 233 Tian He Road North
Unit 1308 – 13th Floor
Guangzhou, China
Tel: +8620 3891 2000
Fax: +8620 3891 2111
E-mail: china-info@brocade.com

Contents

CHAPTER 10
CONFIGURING FWLB FOR LAYER 2 FIREWALLS 10-1

ADDITIONAL FIREWALL CONFIGURATIONS A-1

Chapter 1
About this Guide

This guide describes the features of provides configuration procedures for the Firewall Load Balancing features of the Brocade® ServerIron devices.

Audience

This guide is intended for network engineers with a basic knowledge of switching, routing, and application traffic management.

Conventions

This guide uses the following typographical conventions to describe information:

Italic	Highlights the title of another publication or emphasizes a word or phrase.
`Bold`	Indicates code that is entered exactly as shown.
Bold	Indicates a command or keyword that can be entered exactly as is.

NOTE: A note emphasizes an important fact or calls your attention to a dependency.

WARNING: A warning calls your attention to a possible hazard that can cause injury or death.

CAUTION: A caution calls your attention to a possible hazard that can damage equipment.

Related Documentation

For more information, refer to the following Brocade Communications Systems ServerIron documentation:

- *Release Notes for ServerIron Switch and Router Software TrafficWorks 10.2.01* – provides a list of new features and enhancements, upgrade procedures, and bug fixes.

- *ServerIron TrafficWorks Graphical User Interface* – provides details on the graphical user interface for the ServerIron family of application delivery controllers.
- *ServerIron TrafficWorks Server Load Balancing Guide* – describes basic Server Load Balancing configurations for the ServerIron product family. It covers the following features: Server Load Balancing, Stateless Server Load Balancing, Health Checks, Layer 7 Content Switching, and High Availability
- *ServerIron TrafficWorks Advanced Server Load Balancing Guide* – discusses Advanced Server Load Balancing concepts for the ServerIron product family. It covers the following features: are SIP Server Load Balancing, Transparent Cache Switching, IDS Server Load Balancing, HTTP Compression, and Total Content Analysis
- *ServerIron TrafficWorks Global Server Load Balancing Guide* – explains how one can achieve site level redundancy and data center site failure protection using Global Server Load Balancing feature of ServerIron
- *ServerIron TrafficWorks Security Guide* – describes Security features of ServerIron product family. It covers the following features: are Secure Socket Layer (SSL) Acceleration, Web Application Firewall, Deep Packet Scan, Access Control List, and Network Address Translation
- *ServerIron TrafficWorks Administration Guide* – discusses different administrative configurations for the ServerIron product family.
- *ServerIron TrafficWorks Switching and Routing Guide* – describes switching and routing configurations on the ServerIron product family
- *ServerIron TrafficWorks Firewall Load Balancing Guide* – discusses firewall load balancing designs that are built using ServerIron application controllers
- *Brocade ServerIron Chassis Hardware Installation Guide* – provides the physical characteristics, power consumption, and performance capabilities of the ServerIron chassis switch families, and explains how to set up and install the switches and their modules.
- Brocade *Management Information Base Reference* – presents the Simple Network Management Protocol (SNMP) Management Information Base (MIB) objects that are supported on Brocade devices.

The latest version of these guides are posted at http://www.brocade.com/ethernetproducts.

If you find errors in the guides, send an e-mail to documentation@brocade.com

Getting technical help

To contact Technical Support, go to *http://www.brocade.com/services-support/index.page* for the latest e-mail and telephone contact information.

Document feedback

Quality is our first concern at Brocade and we have made every effort to ensure the accuracy and completeness of this document. However, if you find an error or an omission, or you think that a topic needs further development, we want to hear from you. Forward your feedback to:

documentation@brocade.com

Provide the title and version number of the document and as much detail as possible about your comment, including the topic heading and page number and your suggestions for improvement.

Chapter 2
New Features and Enhancements

This chapter lists new ServerIron features by release, and directs you to their descriptions in the documentation. This chapter contains information about the following releases:

Software Dependencies for Hardware Platforms

- The ServerIron WSM7 management module requires software release 09.4.00l or later.
- 3-slot chassis (GT-C series or SI 350) is supported from software release 09.4.00g onwards.
- ServerIron 4G series is supported from release 09.5.02a onwards.
- The software enhancements/features available on chassis based systems with release 10.0.00a are available on 4G family from software release 10.0.00 onwards.

Chapter 3
ServerIron FWLB Overview

Firewall Load Balancing (FWLB) allows the ServerIron to balance traffic on multiple firewalls. The ServerIron supports the following FWLB topologies: Basic FWLB, High Availability (HA) FWLB, and Multizone FWLB.

NOTE: The ServerIron does not currently support the following topologies: FWLB + NAT, FWLB + Layer 7, FWLB + SYN Proxy.

This chapter contains the following sections:

Understanding ServerIron FWLB

This section contains the following sections:

Firewall Environments

ServerIron supports load balancing across the following firewall environments:

Synchronous Firewall Environments

In general, firewalls that are synchronized allow the in and out traffic of conversations to pass through multiple firewalls. The firewalls exchange information about the conversation so that the inbound or outbound traffic for the conversation does not need to be revalidated each time it tries to use a different firewall. Although the firewalls themselves are synchronized, you will still need to configure paths on the ServerIrons.

Asynchronous Firewall Environments

Asynchronous firewalls do not exchange information about conversations. Traffic must be revalidated each time it arrives at a new firewall. Path information you configure on the ServerIron provides synchronization for the asynchronous firewalls, thus reducing the overhead caused by needless revalidations.

NAT Firewall Environments

Firewalls that perform NAT can translate private network addresses (for example, 10.0.0.1) on the private side of the firewall into Internet addresses (for example, 209.157.22.26) on the public side of the firewall.

Dynamic Route Environments

ServerIrons in IronClad (high-availability) configurations automatically block Layer 3 route traffic at the backup ServerIron to avoid loops, thus simplifying configuration in these environments. See “Router Paths” on page 3-9.

Static Route Environments

Firewalls in static route environments have static or default routes, as do the external (Internet) and internal routers.

Layer 2 Firewall Environments

Layer 2 firewalls do not route (as Layer 3 firewalls do), so the path configuration is slightly different from the path configuration for Layer 3 firewalls.

NOTE: In all types of FWLB configurations, the ServerIrons must be able to reach the firewalls at Layer 2. Thus the firewalls must be directly attached to the ServerIrons or attached to them through Layer 2 devices.

Load Balancing Paths

To send traffic through firewalls, the ServerIron uses paths. A path consists of the following information:

- Path ID

  The path ID is a number that identifies the path. In a basic FWLB configuration, the paths go from one ServerIron to the other through the firewalls. In IronClad FWLB, additional paths go to routers. On each ServerIron, the path IDs must be contiguous (with no gaps), starting with path ID 1.

- ServerIron port

  The number of the port that connects the ServerIron to the firewall.

- Destination IP address

  The management address of the ServerIron or Layer 2 switch on the other side of the firewall. The ServerIron on the private network side and the other ServerIron or Layer 2 switch are the end points of the data path through the firewall. If the path goes to a router, this parameter is the IP address of the firewall’s interface with the ServerIron.

- Next-hop IP address

  The IP address of the firewall interface connected to this ServerIron.

Figure 3.1 shows an example of FWLB paths.

Figure 3.1 Example of FWLB Paths

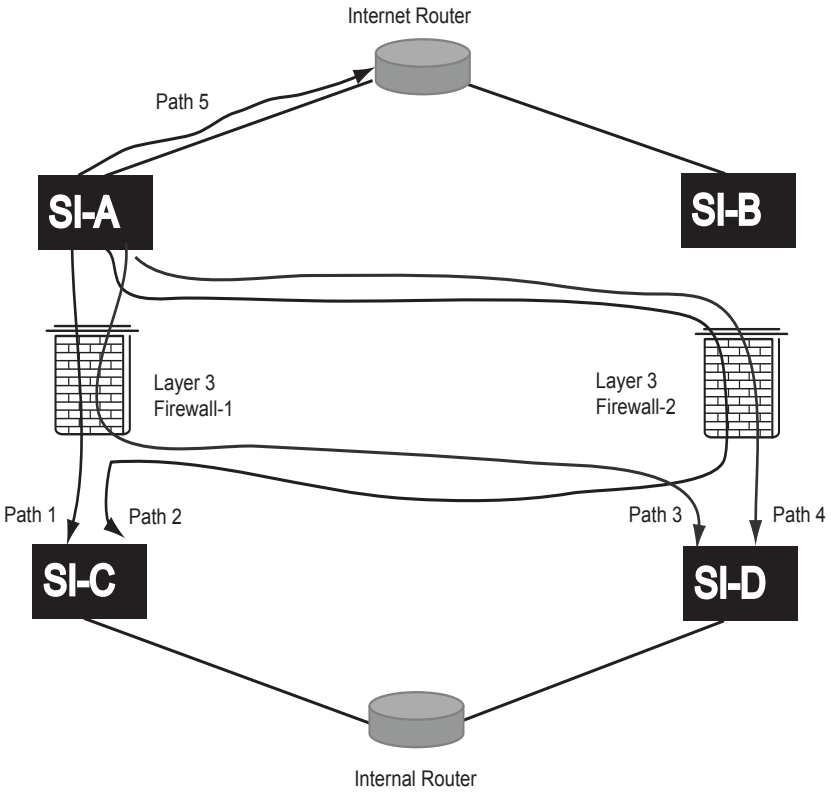

This example above shows the following paths:

- Path 1—ServerIron A through Firewall 1 to ServerIron C
- Path 2—ServerIron A through Firewall 2 to ServerIron C
- Path 3—ServerIron A through Firewall 1 to ServerIron D
- Path 4—ServerIron A through Firewall 2 to ServerIron D
- Path 5—ServerIron A to Internet router.

To ensure proper synchronization of traffic through the firewalls, the paths must be symmetrical. This means that on each ServerIron, the order of next-hop addresses must match. Thus, if you are configuring IronClad FWLB for Layer 3 firewalls, you must configure the paths so that the firewall interfaces are listed in the same order. For example, if the configuration contains four firewalls and you number them 1 – 4 from left to right, the paths on each ServerIron must be configured so that firewalls' next-hop addresses match (the interface for firewall 1 is in the first path, the interface for firewall 2 is in the second path, and so on).

Firewall Selection

Once a ServerIron has selected a firewall for a given traffic flow (source-destination pair of IP addresses), the ServerIron uses the same firewall for subsequent traffic in the same flow. For example, if the ServerIron selects firewall FW1 for the first packet the ServerIron receives with source address 1.1.1.1 and destination address 2.2.2.2, the ServerIron uses FW1 for all packets of flows from 1.1.1.1 to 2.2.2.2.

The ServerIron uses one of the following methods to select a firewall for the first packet:

- Select the firewall based on a hash calculation – used for stateless FWLB
- Select the firewall with the fewest open connections – used for stateful FWLB

Hashing Mechanism

The ServerIrons use the path information along with the hash-mask value for each source-destination pair of IP addresses in the user traffic to consistently send the same source-destination pairs through the same paths. For FWLB, the hash mask must be set to all ones (255.255.255.255 255.255.255.255) to ensure that a given source-destination pair always goes down the same path.

The ServerIron selects a firewall for forwarding a packet based on the packet's hash value (the binary sum of the source and destination addresses). Once the ServerIron assigns a hash value to a given source-destination pair, the ServerIron associates that hash value with a path and always uses the same path for the source-destination pair that has the assigned hash value.

Hashing Based on TCP/UDP Port

You can configure the ServerIron to also hash based on destination TCP or UDP ports. When the ServerIron uses the TCP or UDP port number in addition to the source and destination IP address, traffic with the same source and destination IP address can be load balanced across different paths, based on the destination TCP or UDP port number.

In an IronClad FWLB configuration, you need to configure paths through each of the firewalls to each of the ServerIrons on the other side of the firewalls. You also need to configure a path to the router. You do not configure paths between the ServerIrons in an active-standby pair. These ServerIrons are joined by a dedicated Layer 2 link.

NOTE: The ports in the dedicated link between the active and standby ServerIrons in an IronClad configuration must be in their own port-based VLAN. Add the ports as untagged ports. For added redundancy, configure multiple ports as a trunk group for the dedicated link.

Firewall with Fewest Sessions

FWLB on ServerIron Chassis devices is always stateful. A ServerIron performs ***stateful FWLB*** by creating and using session entries for source and destination traffic flows and associating each flow with a specific firewall.

NOTE: FWLB on the ServerIronXL and ServerIronXL/G is stateless by default and uses the hashing mechanism.

When a ServerIron receives a packet the needs to go through a firewall, the ServerIron checks to see whether it has an existing session entry for the packet.

- If the ServerIron does not have a session entry with the packet's source and destination addresses, the ServerIron creates one. To create the session entry, the ServerIron selects the firewall that has the fewest open sessions with the ServerIron and associates the source and destination addresses of the packet with that firewall.

  The ServerIron also sends the session information to the other ServerIron in the high-availability pair, so that the other ServerIron does not need to create a new session for the same traffic flow.

- If the ServerIron already has a session entry for the packet, the ServerIron forwards the traffic to the firewall in the session entry. All packets with the same source and destination addresses are forwarded to the same firewall. Since the ServerIrons in a high-availability pair exchange session information, the same firewall is used regardless of which ServerIron receives the traffic to be forwarded.

Health Checks

The ServerIron regularly checks the health of the firewall and router paths, and of the applications on the firewalls, if you add applications to the firewall configurations.

Active ServerIrons on each side of a firewall exchange health information for the links in each path by exchanging IP pings through the firewalls. When an active ServerIron on one side of a firewall receives a reply to a ping it sends to the other active ServerIron, on the other side of the firewall, the ServerIron that sent the ping concludes that its partner on the other side of the firewall is operating normally.

The pings are required because a ServerIron can use link-state information to detect when the local link (a link directly attached to a ServerIron port) in a path goes down, but cannot detect when the remote link in the path goes down. If the other ServerIron fails to respond to a ping on a specific port, the ServerIron that sent the ping tries two more times, then determines that the remote link in the path must be down.

NOTE: The health checking mechanism requires that the firewalls be configured to allow ICMP traffic between the two ServerIrons. If the firewalls block the ICMP traffic between ServerIrons, the health check will not work and as a result your IronClad configuration will not function properly.

ServerIrons in an IronClad FWLB configuration also exchange health information. In this case, the ServerIrons exchange packets at Layer 2 and other information related to the link states of the ports that connect the ServerIrons.

In addition to the health checks described above, each ServerIron, whether active or in standby mode, sends IP pings through every path to the other ServerIrons to check the health of the paths. For information about path health checks, see the following section.

Path Health Checks

One of the required FWLB parameters is a separate path from the ServerIron through each firewall to each of the ServerIrons on the other side of the firewall. A path to the ServerIron's gateway router also is required.

By default, the ServerIron performs a Layer 3 health check of each firewall and router path by sending an ICMP ping packet on each path.

- If the ServerIron receives a reply within the allowed amount of time, the ServerIron concludes that the path is good.
- If the ServerIron does not receive a reply within the allowed amount of time, the ServerIron concludes that the path is down.

By default, the ServerIron waits 400 milliseconds for a reply to an ICMP health check packet. If the reply does not arrive, the ServerIron makes two more attempts by default. Therefore, the total amount of time the ServerIron waits for a response is 1.2 seconds by default.

You can increase the total amount of time the ServerIron will wait for a response by increasing the number of attempts. The default maximum number of health check attempts is 3. The valid number of attempts is a value from 3 – 31 on ServerIron Chassis devices or 3 – 31 on other ServerIron models.

Optionally, you can configure the ServerIrons in an FWLB configuration to use Layer 4 TCP or UDP health checks instead of Layer 3 health checks for firewall paths. When you configure a Layer 4 health check, the Layer 3 (ICMP) health check, which is used by default, is disabled. The Layer 4 health check applies only to firewall paths. The ServerIron always uses a Layer 3 (ICMP) health check to test the path to the router.

NOTE: You must configure the same path health check parameters on all the ServerIrons in the FWLB configuration. Otherwise, the paths will not pass the health checks.

Application Health Checks

When you add firewall configuration information to the ServerIron, you also can add information for individual application ports. Adding the application information is optional.

You can specify the following:

- The application's protocol (TCP or UDP) and port number
- The Layer 4 health check state (enabled or disabled) for the application

Adding an application port provides the following benefits:

- The ServerIron includes the source and destination port numbers for the application when it creates session entry. Thus, adding the application port provides more granular load balancing.
- The ServerIron checks the health of the TCP or UDP service used by the application, by sending a Layer 4

TCP or UDP health check to the firewall.

Layer 4 health checks are enabled by default. However, you can disable the Layer 4 health checks globally or on individual application on individual firewalls.

The ServerIron performs the Layer 4 TCP and UDP health checks as follows:

- TCP health check – The ServerIron checks the TCP port's health based on a TCP three-way handshake:
  - The ServerIron sends a TCP SYN packet to the port on the firewall.
  - The ServerIron expects the firewall to respond with a SYN ACK.
  - If the ServerIron receives the SYN ACK, the ServerIron sends a TCP RESET, satisfied that the TCP port is alive.
- UDP health check – The ServerIron sends a UDP packet with garbage (meaningless) data to the UDP port:
  - If the firewall responds with an ICMP "Port Unreachable" message, the ServerIron concludes that the port is not alive.
  - If the server does not respond at all, the ServerIron assumes that the port is alive and received the garbage data. Since UDP is a connectionless protocol, the ServerIron and other clients do not expect replies to data sent to a UDP port. Thus, lack of a response indicates a healthy port.

NOTE: To configure a Layer 4 or Layer 7 application health check, use the procedures in the "Configuring Health Checks" section of the "Configuring Port and Health Check Parameters" chapter in the *Foundry ServerIron Installation and Configuration Guide*. The command syntax and behavior of Layer 4 and Layer 7 health checks is the same regardless of whether you are configuring them for SLB, TCS, or FWLB.

Basic FWLB Topology

You can configure basic FWLB by deploying one ServerIron on the enterprise side of the firewalls and another ServerIron on the Internet side of the firewalls.

A basic FWLB topology uses two ServerIrons to load balance traffic across Layer 3 firewalls. The firewalls can be synchronous or asynchronous.

In the basic configuration, one ServerIron connects to all the firewalls on the private network side. The other ServerIron connects to all the firewalls on the Internet side. The ServerIron(s) balances firewall traffic flows across the firewalls.

Figure 3.2 shows an example of a basic FWLB topology.

Figure 3.2 Basic FWLB Topology

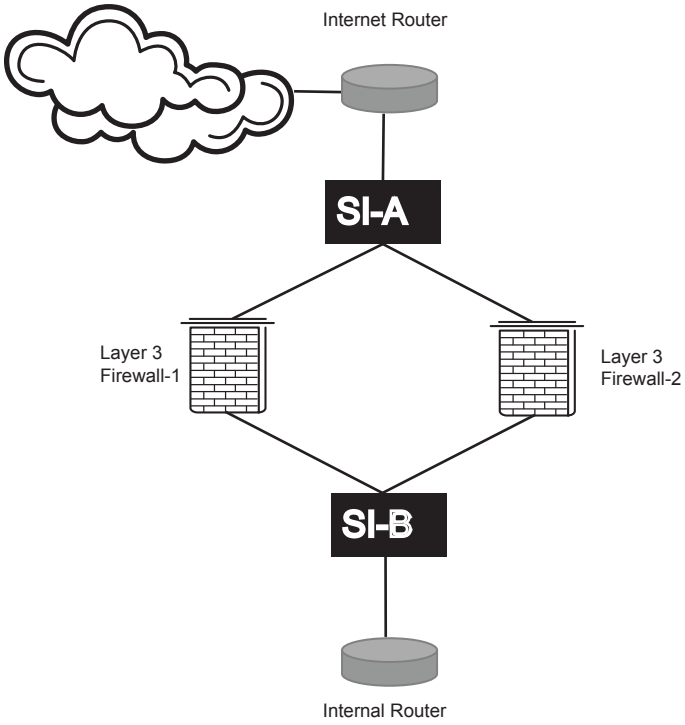

As shown in this example, each ServerIron is configured with paths through the firewalls to the other ServerIron. The ServerIrons use these paths as part of the load balancing mechanism to ensure that traffic for a given IP source and IP destination always passes through the same firewall. All FWLB configurations require paths.

HA FWLB Topology

For high availability (HA), you can deploy pairs of ServerIrons in active-active configurations on each side of the firewalls. In an Active-Active configuration, both ServerIrons in a high-availability pair actively load balance FWLB traffic. Active-Active operation provides redundancy in case a ServerIron becomes unavailable, while enhancing performance by using both ServerIrons to process and forward traffic.

HA FWLB on ServerIron Chassis devices is always stateful. Each ServerIron sends session information about its active traffic flows to the other ServerIron. If a failover occurs, the ServerIron that is still active can provide service for the other ServerIron traffic flows using the session information provided by the other ServerIron.

In an HA topology using ServerIron Chassis devices, both ServerIrons actively load balance traffic to the firewalls. If one of the ServerIrons becomes unavailable, the other ServerIron automatically takes over load balancing for the sessions that were on the unavailable ServerIron.

Figure 3.3 shows an example of HA FWLB.

Figure 3.3 HA FWLB Topology

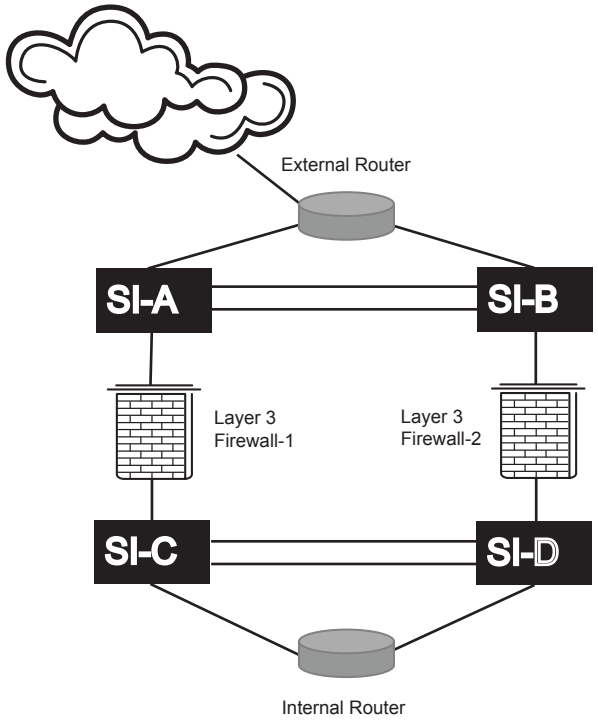

In this example, clients access the application servers on the private network through one of two routers, each of which is connected to a ServerIron. The ServerIrons create session entries for new traffic flows, including assignment of a firewall. The ServerIrons then use the session entries to forward subsequent traffic in the flow to the same firewall.

Failover

In Active-Active FWLB, if one of the ServerIrons becomes unavailable, the other ServerIron takes over for the unavailable ServerIron. The ServerIrons use the following parameters to manage failover:

- ServerIron priority (Active-Standby only) – You can specify a priority from 0 – 255 on each ServerIron. The ServerIron with the higher priority is the default active ServerIron. Specifying the priority is required.

  NOTE: If you specify 0, the CLI removes the priority. When you save the configuration to the startup-config file, the **sym-priority** command is removed. Use this method to remove the priority. You cannot remove the priority using the **no sym-priority** command.

  NOTE: The priority parameter does not apply to Active-Active configurations.

- Path tolerance – Optionally, you also can configure a minimum number of firewall paths and router paths that must be available.

By default, failover occurs if the health checks between the ServerIrons reveal that the active ServerIron has lost a path link. In configurations that contain numerous paths, unstable links can cause frequent failovers, which may be unnecessary and undesirable. To prevent frequent failovers (flapping), you can specify tolerances for the number of good firewall paths and the number of good router paths.

When you configure tolerances, you specify the minimum number of good path links to routers and to firewalls you are requiring the ServerIron to have. So long as the ServerIron has the minimum required number of good links, the ServerIron remains active, even if a link does become unavailable. However, if the number of unavailable links exceeds the minimum requirement you configure and as a result the ServerIron has less available paths than its

active-standby partner, failover to the standby ServerIron occurs. At this point, the standby ServerIron remains active only so long as the number of good paths meets or exceeds the minimums you have configured.

Only if the number of paths is less than the configured minimum and less than the number of available paths on the other ServerIron does failover occur. If the number of paths remains equal on each ServerIron, even if some paths are unavailable on each ServerIron, failover does not occur.

You configure the minimums for firewall paths and router paths separately. The default tolerances are equal to the number of paths of each type you configure. For example, if a ServerIron has four paths through firewalls, the default minimum number of firewall paths required is also four.

Router Paths

IronClad FWLB configurations require paths to the routers in addition to paths to the firewalls. The router paths are required so the ServerIrons can ping the router links to assess their health.

In IronClad FWLB configurations, the standby ServerIrons block Layer 3 OSPF, IGRP, and RIP traffic on the standby paths. This means that the ServerIrons block traffic between routers on different sides of the firewalls if the traffic uses the standby paths. After a failover to a standby ServerIron, the traffic pattern changes. The active ServerIrons allow Layer 3 traffic between routers to pass through the firewalls on the active paths, while blocking the Layer 3 traffic on the standby paths.

NOTE: If you have configured a default route between the routers, the route will work only when the ServerIron through which the route passes is active. If the ServerIron is in standby mode, the route is blocked.

Multizone FWLB Topology

Figure 3.4 shows an example of Multizone Basic FWLB.

Figure 3.4 Multizone Basic FWLB

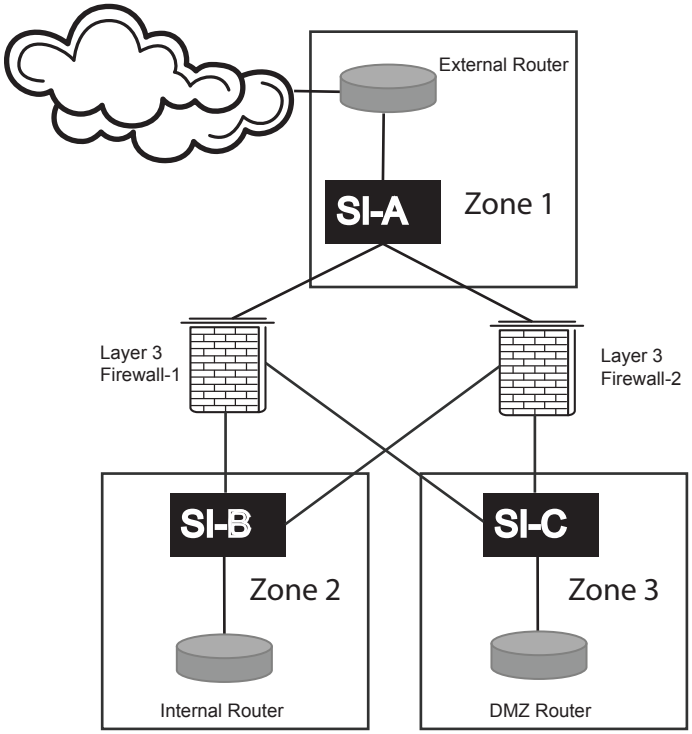

In this example,

Figure 3.5 shows an example of Multizone HA FWLB.

Figure 3.5 Multizone HA FWLB

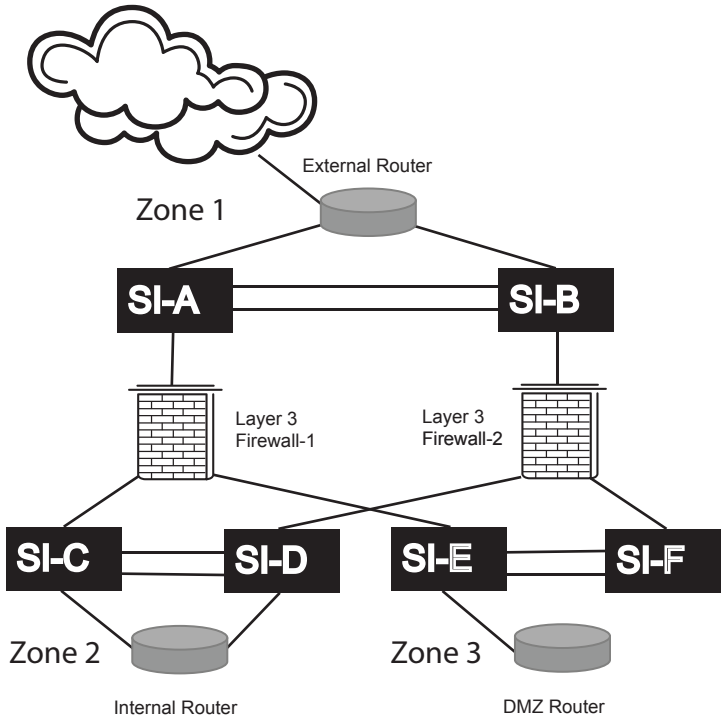

Configuration Guidelines

NOTE: Move the following to the configuration chapter

Use the following guidelines when configuring a ServerIron for FWLB.

- The ServerIron supports one firewall group, group 2. By default, all ServerIron ports belong to this firewall group.
- You must configure a separate path on each ServerIron for each firewall. The paths ensure that firewall traffic with a given pair of source and destination IP addresses flows through the same firewall each time. Thus, the paths reduce firewall overhead by eliminating unnecessary revalidations.

  NOTE: Path configuration is required for all load balancing configurations, whether the firewalls are synchronous or asynchronous.

- You cannot use the features described in the "Configuring Layer 7 Switching" chapter of the *Foundry ServerIron Installation and Configuration Guide* ***and*** FWLB on the same ServerIron.

Configuration Guidelines for FWLB in IronCore Systems

Use the following guidelines to configure FWLB in IronCore systems. Refer to the *ServerIron Chassis L4-7 Configuration Guide* for additional detail and any known limitations.

1. In releases 07.2.xx and 08.x.xx, global firewall policies must be configured for FWLB. Beginning with release 09.3.01, firewall policies are not required. Instead of firewall polices, you must configure the **client-interface** and **server-interface** commands on the interfaces to distribute traffic to WSM CPUs. Refer to the Release Notes for release 09.0.00 or to the *ServerIron Chassis L4-7 Configuration Guide* for more information on these two commands.
2. Rules for configuring external and internal ServerIrons in a FWLB only configuration:
   - The **server-interface** command is required on interfaces connected to firewalls.
   - The **client-interface** command is required on interfaces connected to routers and clients.
   - In high availability configurations, the **server-interface** and the **client-interface** commands *should not be configured* on interfaces used for session synchronization and firewall partner ports.
   - You may connect hosts (which can act as clients or servers) directly to ServerIrons. The **client-interfaces** command must be enabled on all the interfaces connected to these hosts.
3. Rules for External SLB+FWLB configuration, where SLB+FWLB are configured on the external ServerIron and FWLB is configured on the internal ServerIron):
   - When configuring internal and external ServerIrons, follow the rules described in Step 2 above. They apply to both external and internal ServerIrons.
   - Typically, real servers are not attached to external ServerIrons in this configuration. If the configuration requires real servers to be attached to external ServerIrons, the **server-interface** command should be enabled on interfaces connected to servers.
4. Rules for Internal SLB+FWLB configuration, where SLB+FWLB are configured on the internal ServerIron and FWLB on the external ServerIron):
   - On the external ServerIrons, follow the rules described in Step 2 above.
   - On the internal ServerIrons, the **client-interfaces** command should be enabled on interfaces connected to firewalls. The **server-interfaces** command should be enabled on all the interfaces connected to real servers.
   - If there are any remote servers, the **server-interfaces** command should be enabled on the interface connected to next hop router.
   - In high availability configurations, both the **client-interface** and the **server-interface** command *should not be configured* on interfaces used for session synchronization and firewall partner ports.
   - The **client-interfaces** command should be enabled on Interfaces connected to clients that are directly attached to internal ServerIrons.

Configuration Guidelines for FWLB in JetCore system

Beginning with this release, firewall policies are not required for FWLB configuration. Also, you do not need to configure the **client-interface** and **server-interface** commands. Traffic will be distributed to WSM CPUs according to the Layer4-7 CAM entries created for FWLB and SLB.

FWLB Configuration Limits

Table 3.1 contains the FWLB configuration limits supported by the ServerIron.

Table 3.1: FWLB Configuration Limits

Maximum Firewall Groups	Maximum Firewalls	Maximum Paths	Maximum Zones	Maximum Router Paths
1 (group 2)	16	32	3 (internal, external, dmz)	4

Chapter 4
Configuring Basic FWLB

This chapter describes how to implement commonly used configurations for the following:

- Basic FWLB (configuration without ServerIron redundancy)
- IronClad (active-standby configuration with ServerIron redundancy)

Configuring Basic Layer 3 FWLB

Basic FWLB uses a single ServerIron on the enterprise side of the load balanced firewalls and another ServerIron on the Internet side. Figure 3.2 on page 3-7 shows an example of this type of configuration.

Configuring Basic Layer 3 FWLB

To configure basic Layer 3 FWLB, perform the following tasks.

Table 4.1: Configuration tasks – Basic FWLB

Task	See page...
Configure Global Parameters	
Globally enable FWLB	4-1
Configure Firewall Parameters	
Define the firewalls and add them to the firewall group	4-2
Configure Firewall Group Parameters	
Configure the paths and add static MAC entries for the firewall interfaces with the ServerIron	4-3

Enabling FWLB

To enable FWLB, you configure global IP policies. FWLB for TCP and UDP is controlled independently, so you need to configure a separate global IP policy for each type of traffic.

When you enable FWLB for TCP or UDP globally, all ports that are in the firewall group are enabled for FWLB. All ServerIron ports are in firewall group 2 by default. Thus, if you enable FWLB globally, it affects all physical ports unless you remove ports from firewall groups.

NOTE: The user interface allows you to enable FWLB locally instead of globally. However, local policies are not applicable to FWLB. Enable the feature globally.

To enable FWLB globally, use the following method.

USING THE CLI

Enter the following commands at the global CONFIG level to enable FWLB for all TCP and UDP traffic:

```
ServerIron(config)# ip policy 1 fw tcp 0 global
ServerIron(config)# ip policy 2 fw udp 0 global
```

Syntax: [no] ip policy <policy-num> fw tcp | udp 0 global

The <policy-num> value identifies the policy and can be a number from 1 – 64.

Each policy affects TCP or UDP traffic, so you must specify **tcp** or **udp**.

The value 0 following the **tcp | udp** parameter specifies that the policy applies to all ports of the specified type (TCP or UDP). In this command, "0" is equivalent to "any port number". For FWLB, you must specify "0".

NOTE: Generally, the firewall itself performs validation and authentication for the traffic, so allowing the ServerIron to pass all traffic of the specified type (TCP or UDP) to the firewall simplifies configuration.

Defining the Firewalls and Adding them to the Firewall Group

When FWLB is enabled, all the ServerIron ports are in firewall group 2 by default. However, you need to add an entry for each firewall, then add the firewalls to the firewall group. To add an entry for a firewall, specify the firewall name and IP address. You can specify a name up to 32 characters long.

To define the firewalls shown in Figure 3.2 on page 3-7 and add them to firewall group 2, use the following method.

USING THE CLI

To define the firewalls using the CLI, enter the following commands.

Commands for ServerIron A (External)

```
ServerIron(config)# server fw-name FW1-IPin 209.157.22.3
ServerIron(config-rs-FW1-IPin)# exit
ServerIron(config)# server fw-name FW2-IPin 209.157.22.4
ServerIron(config-rs-FW2-IPin)# exit
ServerIron(config)# server fw-group 2
ServerIron(config-tc-2)# fw-name FW1-IPin
ServerIron(config-tc-2)# fw-name FW2-IPin
```

Commands for ServerIron B (Internal)

```
ServerIron(config)# server fw-name FW1-IPout 209.157.23.1
ServerIron(config-rs-FW1-IPout)# exit
ServerIron(config)# server fw-name FW2-IPout 209.157.23.2
ServerIron(config-rs-FW2-IPout)# exit
ServerIron(config)# server fw-group 2
ServerIron(config-tc-2)# fw-name FW1-IPout
ServerIron(config-tc-2)# fw-name FW2-IPout
```

Syntax: [no] server fw-name <string> <ip-addr>

NOTE: When you add a firewall name, the CLI level changes to the Firewall level. This level is used when you are configuring stateful FWLB.

Syntax: server fw-group 2

This command changes the CLI to firewall group configuration level. The firewall group number is 2. Only one firewall group is supported.

Syntax: [no] fw-name <string>

Adds a configured firewall to the firewall group.

Configuring the Paths and Adding Static MAC Entries

A path is configuration information the ServerIron uses to ensure that a given source and destination IP pair is always authenticated by the same Layer 3 firewall.

Each path consists of the following parameters:

- The path ID – A number that identifies the path. The paths go from one ServerIron to the other through the firewalls. On each ServerIron, the sequence of path IDs must be contiguous (with no gaps), starting with path ID 1. For example, path sequence 1, 2, 3, 4, 5 is valid. Path sequence 1, 3, 5 or 5, 4, 3, 2, 1 is not valid.
- The ServerIron port – The number of the port that connects the ServerIron to the firewall. If your configuration does not require static MAC entries, you can specify a dynamic port (65535) instead of the physical port number for firewall paths. Specifying the dynamic port allows the ServerIron to select the physical port for the path so you don't need to.
- The other ServerIron's or Layer 2 switch's IP address – The management address of the ServerIron or Layer 2 switch on the other side of the firewall. The ServerIron on the private network side and the other ServerIron or Layer 2 switch are the end points of the data path through the firewall.
- The next-hop IP address – The IP address of the firewall interface connected to this ServerIron.

For each type of firewall (Layer 3 synchronous and asynchronous, with or without NAT), you must configure paths between the ServerIrons through the firewalls.

In addition to configuring the paths, you need to create a static MAC entry for each firewall interface attached to the ServerIron.

NOTE: When defining a firewall router path on a port, make sure the port is a **server router-port**.

NOTE: FWLB paths must be fully meshed. When you configure a FWLB path on a ServerIron, make sure you also configure a reciprocal path on the ServerIron attached to the other end of the firewalls. For example, if you configure four paths to four separate firewalls, make sure you configure four paths on the other ServerIron.

NOTE: For many configurations, static MAC entries are required. Where required, you must add a static MAC entry for each firewall interface with the ServerIron. The FWLB configuration examples in this guide indicate whether static MAC entries are required.

To configure a path and add static MAC entries, use one of the following methods.

USING THE CLI

To configure the paths and static MAC entries for the configuration shown in Figure 3.2 on page 3-7, enter the following commands. Enter the first group of commands on ServerIron A. Enter the second group of commands on ServerIron B.

Commands for ServerIron A (External)

```
ServerIron(config)# server fw-group 2
ServerIron(config-tc-2)# fwall-info 1 3 209.157.23.3 209.157.22.3
ServerIron(config-tc-2)# fwall-info 2 5 209.157.23.3 209.157.22.4
ServerIron(config-tc-2)# exit
ServerIron(config)# static-mac-address abcd.4321.34e0 ethernet 3 high-priority
router-type
ServerIron(config)# static-mac-address abcd.4321.34e1 ethernet 5 high-priority
router-type
```



```
ServerIron(config)# write mem
```

Commands for ServerIron B (Internal)

```
ServerIron(config)# server fw-group 2
ServerIron(config-tc-2)# fwall-info 1 1 209.157.22.2 209.157.23.1
ServerIron(config-tc-2)# fwall-info 2 2 209.157.22.2 209.157.23.2
ServerIron(config-tc-2)# exit
ServerIron(config)# static-mac-address abcd.4321.34e2 ethernet 1 high-priority
router-type
ServerIron(config)# static-mac-address abcd.4321.34e3 ethernet 2 high-priority
router-type
ServerIron(config)# write mem
```

Command Syntax

Syntax: server fw-group 2

Syntax: [no] fwall-info <path-num> <portnum> <other-ServerIron-ip> <next-hop-ip>

The syntax for adding static MAC entries differs depending on whether you are using a stackable or chassis ServerIron.

Syntax for chassis devices:

Syntax: [no] static-mac-address <mac-addr> ethernet <portnum> [priority <0-7>] [host-type | router-type]

Syntax for stackable devices:

Syntax: static-mac-address <mac-addr> ethernet <portnum> [to <portnum> ethernet <portnum>] [normal-priority | high-priority] [host-type | router-type | fixed-host]

The priority can be 0 – 7 (0 is lowest and 7 is highest) for chassis devices and either normal-priority or high-priority for stackable devices.

The defaults are **host-type** and **0** or **normal-priority**.

NOTE: The static MAC entries are required. You must add a static MAC entry for each firewall interface with the ServerIron. In addition, you must use the **high-priority** and **router-type** parameters with the **static-mac-address** command. These parameters enable the ServerIron to use the address for FWLB.

NOTE: If you enter the command at the global CONFIG level, the static MAC entry applies to the default port-based VLAN (VLAN 1). If you enter the command at the configuration level for a specific port-based VLAN, the entry applies to that VLAN and not to the default VLAN.

Configuration Example for Basic Layer 3 FWLB

The following sections show all the ServerIron commands you would enter on each ServerIron to implement the configuration shown in Figure 3.2 on page 3-7.

Commands on ServerIron A (External)

Enter the following commands to configure FWLB on ServerIron A.

```
ServerIronA(config)# server fw-name FW1-IPin 209.157.22.3
ServerIronA(config-rs-FW1-IPin)# exit
ServerIronA(config)# server fw-name FW2-IPin 209.157.22.4
ServerIronA(config-rs-FW2-IPin)# exit
```

The commands above add two firewalls, FW1-IPin and FW2-IPin.

The following commands configure parameters for firewall group 2. The **fwall-info** commands configure the paths for the firewall traffic. Each path consists of a path ID, the ServerIron port attached to the firewall, the IP address of the ServerIron at the other end of the path, and the next-hop IP address (usually the firewall interface connected to this ServerIron). Make sure you configure reciprocal paths on the other ServerIron, as shown in the section containing the CLI commands for ServerIron B.

NOTE: Path information is required even if the firewalls are synchronized.

The **fw-name** <firewall-name> command adds the firewalls to the firewall group.

```
ServerIronA(config)# server fw-group 2
ServerIronA(config-tc-2)# fw-name FW1-IPin
ServerIronA(config-tc-2)# fw-name FW2-IPin
ServerIronA(config-tc-2)# fwall-info 1 3 209.157.23.3 209.157.22.3
ServerIronA(config-tc-2)# fwall-info 2 5 209.157.23.3 209.157.22.4
ServerIronA(config-tc-2)# exit
```

The following commands add static MAC entries for the MAC addresses of the firewall interfaces connected to the ServerIron. Notice that the QoS priority is configured as **high-priority** and the **router-type** parameter is specified. These parameters are required. You must specify **high-priority** and **router-type**.

NOTE: To ensure proper operation, always configure the path IDs so that the IDs consistently range from lowest path ID to highest path ID for the firewalls. For example, in Figure 3.2 on page 3-7, the path IDs should range from lowest to highest beginning with the firewall interface at the upper left of the figure.

To ensure smooth operation, you might want to depict your firewalls in a vertical hierarchy as in Figure 3.2 on page 3-7, label the interfaces with their IP addresses, then configure the paths so that the path IDs to the interfaces range from lowest to highest path ID starting from the uppermost firewall interface.

```
ServerIronA(config)# static-mac-address abcd.4321.34e0 ethernet 3 high-priority
router-type
ServerIronA(config)# static-mac-address abcd.4321.34e1 ethernet 5 high-priority
router-type
```

The following commands configure global policies to enable FWLB. Global or local policies are required for FWLB. The first **ip policy** command in this example configures the ServerIron to perform FWLB for all TCP traffic. The value “0” is equivalent to “any” and means the ServerIron should perform FWLB for all TCP traffic. The second **ip policy** command enables FWLB for all UDP traffic.

```
ServerIronA(config)# ip policy 1 fw tcp 0 global
ServerIronA(config)# ip policy 2 fw udp 0 global
ServerIronA(config)# write memory
```

Commands on ServerIron B (Internal)

Enter the following commands to configure FWLB on ServerIron B. Notice that the **fwall-info** commands configure paths that are reciprocal to the paths configured on ServerIron A. Path 1 on each ServerIron goes through one of the firewalls while path 2 goes through the other firewall.

```
ServerIronB(config)# server fw-name FW1-IPout 209.157.23.1
ServerIronB(config-rs-FW1-IPout)# exit
ServerIronB(config)# server fw-name FW2-IPout 209.157.23.2
ServerIronB(config-rs-FW2-IPout)# exit
ServerIronB(config)# server fw-group 2
ServerIronB(config-tc-2)# fw-name FW1-IPout
ServerIronB(config-tc-2)# fw-name FW2-IPout
ServerIronB(config-tc-2)# fwall-info 1 1 209.157.22.2 209.157.23.1
ServerIronB(config-tc-2)# fwall-info 2 2 209.157.22.2 209.157.23.2
ServerIronB(config-tc-2)# exit
ServerIronB(config)# static-mac-address abcd.4321.34e2 ethernet 1 high-priority
router-type
```



```
ServerIronB(config)# static-mac-address abcd.4321.34e3 ethernet 2 high-priority
router-type
ServerIronB(config)# ip policy 1 fw tcp 0 global
ServerIronB(config)# ip policy 2 fw udp 0 global
ServerIronB(config)# write memory
```

Configuration Examples with Layer 3 Routing Support

NOTE: Layer 3 routing is supported only on ServerIron Chassis devices running software release 08.0.00 or later.

This section shows examples of commonly used ServerIron basic FWLB deployments with Layer 3 configurations. The ServerIrons in these examples perform Layer 3 routing in addition to Layer 2 and Layer 4 – 7 switching.

Generally, the steps for configuring Layer 4 – 7 features on a ServerIron running Layer 3 are similar to the steps on a ServerIron that is not running Layer 3. The examples focus on the Layer 3 aspects of the configurations.

This section contains the following configuration examples:

- "Basic FWLB with One Sub-Net and One Virtual Routing Interface" on page 4-6
- "Basic FWLB with Multiple Sub-Nets and Multiple Virtual Routing Interfaces" on page 4-9

NOTE: The basic FWLB configurations shown in these examples are the ones that are supported. If you need to use the ServerIron's Layer 3 routing support in a FWLB configuration that is not shown, contact Brocade Communications Systems.

Basic FWLB with One Sub-Net and One Virtual Routing Interface

Figure 4.1 shows an example of a basic FWLB configuration in which each ServerIron is in only one sub-net. On each ServerIron, a virtual routing interface is configured on all the ports in VLAN 1 (the default VLAN), and an IP sub-net address is configured on the virtual routing interface.

The ServerIron supports dynamic routing protocols, including RIP and OSPF. However, some firewalls do not support dynamic routing and instead require static routes. The network in this example assumes that the firewalls do not support dynamic routing. Since the network uses static routes, each ServerIron is configured with an IP default route that uses one of the firewall interfaces as the next hop for the route.

In addition, the client and server in this network each use a firewall interface as the default gateway. When this is the case, you need to do one of the following:

- Configure each ServerIron with static MAC entries for the firewall interfaces. This example uses the static entries.
- Configure the clients and servers to use the ServerIron itself as the default gateway.

Figure 4.1 Basic FWLB in One Subnet

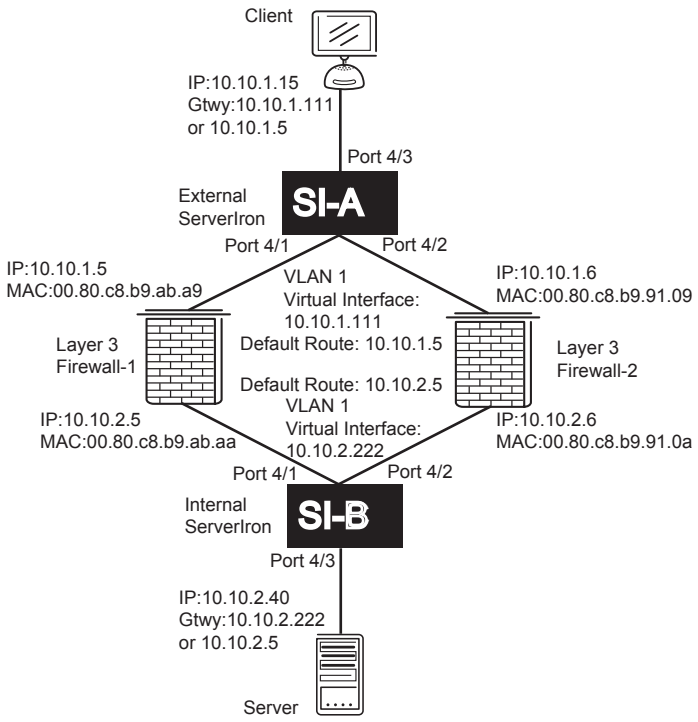

The following sections show the CLI commands for configuring the basic FWLB implementation in Figure 4.1.

Commands on the External ServerIron

The following commands change the CLI to the global CONFIG level, then change the hostname to "SI-External".

```
ServerIron> enable
ServerIron# configure terminal
ServerIron(config)# hostname SI-External
```

The following commands configure a virtual routing interface on VLAN 1 (the default VLAN), then configure an IP address on the interface. The virtual routing interface is associated with all the ports in the VLAN. In this case, since all the ServerIron ports are in the default VLAN, the virtual routing interface is associated with all the ports on the device.

```
SI-External(config)# vlan 1
SI-External(config-vlan-1)# router-interface ve 1
SI-External(config-vlan-1)# exit
SI-External(config)# interface ve 1
SI-External(config-ve-1)# ip address 10.10.1.111 255.255.255.0
SI-External(config-ve-1)# exit
```

The following command configures an IP default route. The first two "0.0.0.0" portions of the address are the IP address and network mask. Always specify zeroes when configuring an IP default route. The third value is the IP address of the next-hop gateway for the default route. In most cases, you can specify the IP address of one of the firewalls as the next hop. Specifying the default route is the Layer 3 equivalent of specifying the default gateway.

```
SI-External(config)# ip route 0.0.0.0 0.0.0.0 10.10.1.5
```

The following commands add the firewall definitions. In this example, port HTTP is configured on each firewall. Specifying the application ports on the firewalls is optional. If you configure an application port on a firewall, load balancing is performed for the configured port. All traffic from a given client for ports that are not configured is sent to the same firewall.

```
SI-External(config)# server fw-name fw1 10.10.1.5
SI-External(config-rs-fw1)# port http
```



```
SI-External(config-rs-fw1)# exit
SI-External(config)# server fw-name fw2 10.10.1.6
SI-External(config-rs-fw2)# port http
SI-External(config-rs-fw2)# exit
```

The following commands add the firewall definitions to the firewall port group (always group 2). The firewall group contains all the ports in VLAN 1 (the default VLAN).

```
SI-External(config)# server fw-group 2
SI-External(config-tc-2)# fw-name fw1
SI-External(config-tc-2)# fw-name fw2
```

The following commands add the paths through the firewalls to the other ServerIron. Each path consists of a path number, a ServerIron port number, the IP address at the other end of the path, and the next-hop IP address. In this example, the topology does not contain routers other than the ServerIrons. If your topology does contain other routers, configure firewall paths for the routers too. For router paths, use the same IP address as the path destination and the next hop.

NOTE: The path IDs must be in contiguous, ascending numerical order, starting with 1. For example, path sequence 1, 2, 3 , 4 is valid. Path sequence 4, 3, 2, 1 or 1, 3, 4, 5 is not valid.

```
SI-External(config-tc-2)# fwall-info 1 4/1 10.10.2.222 10.10.1.5
SI-External(config-tc-2)# fwall-info 2 4/2 10.10.2.222 10.10.1.6
```

The following command sets the load balancing method to balance requests based on the firewall that has the least number of connections for the requested service. Since the firewall definitions above specify the HTTP service, the ServerIron will load balance requests based on the firewall that has fewer HTTP session entries in the ServerIron session table.

```
SI-External(config-tc-2)# fw-predictor per-service-least-conn
SI-External(config)# exit
```

The following commands add static MAC entries for the firewall interfaces with the ServerIron. The static MAC entries are required only if the configuration uses static routes and a single virtual routing interface, as in this example, and if the default gateway for the client or server is the firewall. If the configuration uses a dynamic routing protocol (for example, RIP or OSPF), the static entries are not required. Alternatively, the static entries are not required if you use the ServerIron itself as the default gateway for the client or the server. For example, the static entries are not required if you configure the client to use 10.10.1.111 as its default gateway.

```
SI-External(config)# vlan 1
SI-External(config-vlan-1)# static-mac-address 0080.c8b9.aba9 ethernet 4/1
priority 1 router-type
SI-External(config-vlan-1)# static-mac-address 0080.c8b9.9109 ethernet 4/2
priority 1 router-type
SI-External(config-vlan-1)# exit
```

The following commands assign FWLB processing for all forwarding modules to the same WSM CPU. The device uses the same CPU to process all FWLB traffic. You must assign all the traffic to the same WSM CPU. The commands in this example assign traffic on the forwarding modules in slots 3 and 4 to WSM CPU 1 on the Web Switching Management Module in slot 2.

```
SI-External(config)# wsm wsm-map slot 3 wsm-slot 2 wsm-cpu 1
SI-External(config)# wsm wsm-map slot 4 wsm-slot 2 wsm-cpu 1
```

NOTE: For simplicity, the configuration of the other ServerIron in this example does not include **wsm wsm-map** commands. However, the commands you need to enter depend on the slot locations of the modules in the device and the WSM CPU you want to use.

The following commands enable FWLB.

```
SI-External(config)# ip l4-policy 1 fw tcp 0 global
SI-External(config)# ip l4-policy 2 fw udp 0 global
```

The following command saves the configuration changes to the startup-config file.

```
SI-External(config)# write memory
```


Commands on the Internal ServerIron

```
ServerIron> enable
ServerIron# configure terminal
ServerIron(config)# hostname SI-Internal
SI-Internal(config)# vlan 1
SI-Internal(config-vlan-1)# router-interface ve 1
SI-Internal(config-vlan-1)# exit
SI-Internal(config)# interface ve 1
SI-Internal(config-ve-1)# ip address 10.10.2.222 255.255.255.0
SI-Internal(config-ve-1)# exit
SI-Internal(config)# ip route 0.0.0.0 0.0.0.0 10.10.2.5
SI-Internal(config)# server fw-name fw1 10.10.2.5
SI-Internal(config-rs-fw1)# port http
SI-Internal(config-rs-fw1)# exit
SI-Internal(config)# server fw-name fw2 10.10.2.6
SI-Internal(config-rs-fw2)# port http
SI-Internal(config-rs-fw2)# exit
SI-Internal(config)# server fw-group 2
SI-Internal(config-tc-2)# fw-name fw1
SI-Internal(config-tc-2)# fw-name fw2
SI-Internal(config-tc-2)# fwall-info 1 4/1 10.10.1.111 10.10.2.5
SI-Internal(config-tc-2)# fwall-info 2 4/2 10.10.1.111 10.10.2.6
SI-Internal(config-tc-2)# fw-predictor per-service-least-conn
SI-Internal(config)# exit
SI-Internal(config)# vlan 1
SI-Internal(config-vlan-1)# static-mac-address 0080.c8b9.abaa ethernet 4/1
priority 1 router-type
SI-Internal(config-vlan-1)# static-mac-address 0080.c8b9.910a ethernet 4/2
priority 1 router-type
SI-Internal(config-vlan-1)# exit
SI-Internal(config)# ip l4-policy 1 fw tcp 0 global
SI-Internal(config)# ip l4-policy 2 fw udp 0 global
SI-Internal(config)# write memory
```

Basic FWLB with Multiple Sub-Nets and Multiple Virtual Routing Interfaces

Figure 4.2 shows an example of a basic FWLB configuration in which multiple IP sub-net interfaces are configured on each ServerIron. On each ServerIron, the client or server is in one sub-net and the firewalls are in another sub-net. The ports connected to the firewalls are configured in a separate port-based VLAN. The ServerIron's IP interface to the firewalls is configured on a virtual routing interface associated with the ports in the VLAN.

The client and server in this example are each configured to use their locally attached ServerIron as the default gateway, instead of using a firewall interface. Therefore, you do not need to configure static MAC entries for the firewalls on the ServerIron.

Figure 4.2 Basic FWLB in Multiple Sub-nets Using Multiple Routing Interfaces

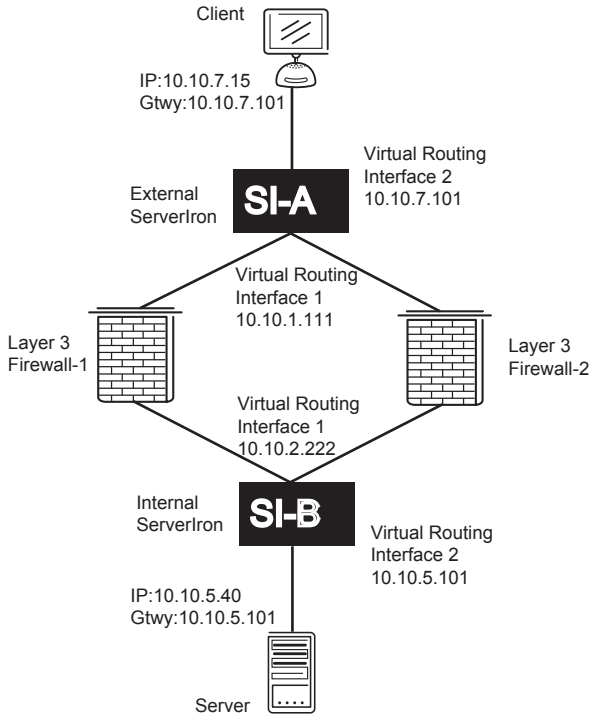

Commands on the External ServerIron

The following commands change the CLI to the global CONFIG level, then change the hostname to "SI-External".

```
ServerIron> enable
ServerIron# configure terminal
ServerIron(config)# hostname SI-External
```

The following commands configure virtual routing interface 1, which is connected to the firewalls. Since both firewalls are in the same sub-net, you must configure the ServerIron's IP interface with the firewalls on a virtual routing interface. Otherwise, you cannot configure the same address on more than port.

The first three commands configure the VLAN. The last two commands configure an IP address on the interface. The IP address is assigned to all the ports in the VLAN associated with the virtual routing interface.

```
SI-External(config)# vlan 10
SI-External(config-vlan-10)# untagged ethernet 4/1 to 4/4
SI-External(config-vlan-10)# router-interface ve 1
SI-External(config-vlan-10)# exit
SI-External(config)# interface ve 1
SI-External(config-ve-1)# ip address 10.10.1.111 255.255.255.0
SI-External(config-ve-1)# exit
```

The following commands configure virtual routing interface 2, which is connected to the client.

```
SI-External(config)# vlan 20
SI-External(config-vlan-20)# untagged ethernet 4/5 to 4/24
SI-External(config-vlan-20)# router-interface ve 2
SI-External(config-vlan-20)# exit
SI-External(config)# interface ve 2
SI-External(config-ve-2)# ip address 10.10.7.101 255.255.255.0
SI-External(config-ve-2)# exit
```

Since Figure 4.2 on page 4-10 shows only one port connected to one client, you could configure the IP address on the physical port attached to the client instead of configuring the address on a separate VLAN. This example uses a virtual routing interface to demonstrate that you can use multiple virtual routing interfaces in your configuration.

The following command configures an IP default route. The first two "0.0.0.0" portions of the address are the IP address and network mask. Always specify zeroes when configuring an IP default route. The third value is the IP address of the next-hop gateway for the default route. In most cases, you can specify the IP address of one of the firewalls as the next hop. Specifying the default route is the Layer 3 equivalent of specifying the default gateway.

```
SI-External(config)# ip route 0.0.0.0 0.0.0.0 10.10.1.5
```

The following commands add the firewall definitions.

```
SI-External(config)# server fw-name fw1 10.10.1.5
SI-External(config-rs-fw1)# port http
SI-External(config-rs-fw1)# exit
SI-External(config)# server fw-name fw2 10.10.1.6
SI-External(config-rs-fw2)# port http
SI-External(config-rs-fw2)# exit
```

The following commands add the firewall definitions to the firewall port group.

```
SI-External(config)# server fw-group 2
SI-External(config-tc-2)# fw-name fw1
SI-External(config-tc-2)# fw-name fw2
```

The following commands add the paths through the firewalls to the other ServerIron. Each path consists of a path number, a ServerIron port number, the IP address at the other end of the path, and the next-hop IP address. In this example, the topology does not contain routers other than the ServerIrons. If your topology does contain other routers, configure firewall paths for the routers too. For router paths, use the same IP address as the path destination and the next hop.

NOTE: The path IDs must be in contiguous, ascending numerical order, starting with 1. For example, path sequence 1, 2, 3 , 4 is valid. Path sequence 4, 3, 2, 1 or 1, 3, 4, 5 is not valid.

```
SI-External(config-tc-2)# fwall-info 1 4/1 10.10.2.222 10.10.1.5
SI-External(config-tc-2)# fwall-info 2 4/2 10.10.2.222 10.10.1.6
```

The following command sets the load balancing method to balance requests based on the firewall that has the least number of connections for the requested service.

```
SI-External(config-tc-2)# fw-predictor per-service-least-conn
SI-External(config-tc-2)# exit
```

The following commands assign FWLB processing for all forwarding modules to the same WSM CPU. The device uses the same CPU to process all FWLB traffic. You must assign all the traffic to the same WSM CPU. The commands in this example assign traffic on the forwarding modules in slots 3 and 4 to WSM CPU 1 on the Web Switching Management Module in slot 2.

```
SI-External(config)# wsm wsm-map slot 3 wsm-slot 2 wsm-cpu 1
SI-External(config)# wsm wsm-map slot 4 wsm-slot 2 wsm-cpu 1
```

NOTE: For simplicity, the configuration of the other ServerIron in this example does not include **wsm wsm-map** commands. However, the commands you need to enter depend on the slot locations of the modules in the device and the WSM CPU you want to use.

The following commands enable FWLB.

```
SI-External(config)# ip l4-policy 1 fw tcp 0 global
SI-External(config)# ip l4-policy 2 fw udp 0 global
```

The following command saves the configuration changes to the startup-config file.

```
SI-External(config)# write memory
```

Commands on the Internal ServerIron

```
ServerIron> enable
ServerIron# configure terminal
ServerIron(config)# hostname SI-Internal
SI-Internal(config)# vlan 10
```



```
SI-Internal(config-vlan-10)# untagged ethernet 4/1 to 4/4
SI-Internal(config-vlan-10)# router-interface ve 1
SI-Internal(config-vlan-10)# exit
SI-Internal(config)# interface ve 1
SI-Internal(config-ve-1)# ip address 10.10.2.222 255.255.255.0
SI-Internal(config-ve-1)# exit
SI-Internal(config)# vlan 20
SI-Internal(config-vlan-20)# untagged ethernet 4/5 to 4/24
SI-Internal(config-vlan-20)# router-interface ve 2
SI-Internal(config-vlan-20)# exit
SI-Internal(config)# interface ve 2
SI-Internal(config-ve-2)# ip address 10.10.5.101 255.255.255.0
SI-Internal(config-ve-2)# exit
SI-Internal(config)# ip route 0.0.0.0 0.0.0.0 10.10.2.5
SI-Internal(config)# server fw-name fw1 10.10.2.5
SI-Internal(config-rs-fw1)# port http
SI-Internal(config-rs-fw1)# exit
SI-Internal(config)# server fw-name fw2 10.10.2.6
SI-Internal(config-rs-fw2)# port http
SI-Internal(config-rs-fw2)# exit
SI-Internal(config)# server fw-group 2
SI-Internal(config-tc-2)# fw-name fw1
SI-Internal(config-tc-2)# fw-name fw2
SI-Internal(config-tc-2)# fwall-info 1 4/1 10.10.1.111 10.10.2.5
SI-Internal(config-tc-2)# fwall-info 2 4/2 10.10.1.111 10.10.2.6
SI-Internal(config-tc-2)# fw-predictor per-service-least-conn
SI-Internal(config-tc-2)# exit
SI-Internal(config)# ip l4-policy 1 fw tcp 0 global
SI-Internal(config)# ip l4-policy 2 fw udp 0 global
SI-Internal(config)# write memory
```


Chapter 5
Configuring HA FWLB

High Availability (HA) FWLB allows the ServerIron Chassis device to actively load balance traffic and provide enhanced performance. This chapter contains the following sections:

Understanding ServerIron FWLB

This section contains the following sections:

Stateful FWLB

FWLB on ServerIron Chassis devices is always stateful. A ServerIron performs ***stateful FWLB*** by creating and using session entries for source and destination traffic flows and associating each flow with a specific firewall.

When a ServerIron receives a packet that needs to go through a firewall, the ServerIron checks to see whether it has an existing session entry for the packet.

- If the ServerIron does not have a session entry with the packet's source and destination addresses, the ServerIron creates one. To create the session entry, the ServerIron selects the firewall that has the fewest open sessions with the ServerIron and associates the source and destination addresses of the packet with that firewall.

  The ServerIron also sends the session information to the other ServerIron in the high-availability pair, so that

the other ServerIron does not need to create a new session for the same traffic flow.

- If the ServerIron already has a session entry for the packet, the ServerIron forwards the traffic to the firewall in the session entry. All packets with the same source and destination addresses are forwarded to the same firewall. Since the ServerIrons in a high-availability pair exchange session information, the same firewall is used regardless of which ServerIron receives the traffic to be forwarded.

Layer 3/4 Sessions

The source and destination addresses in a session entry are Layer 3 or Layer 4.

- A Layer 3 session contains source and destination IP addresses.
- A Layer 4 session entry contains source and destination TCP and UDP port numbers in addition to IP addresses.

The session entry type depends on whether you configure add application ports (TCP or UDP ports) to the firewall configuration information on the ServerIron.

- If you do not configure application ports on a firewall, the ServerIron creates session entries using the source and destination IP addresses only. All packets for a given pair of source and destination IP addresses is always sent to the same firewall.
- If you configure an application port on a firewall, the ServerIron includes the source and destination TCP or UDP port numbers in the session entries for the application. Packets for the same set of source and destination IP addresses can be sent to different firewalls, depending on the source and destination TCP or UDP port numbers in the packets. For example, if you configure TCP port 80 on the firewalls, the ServerIron uses IP addresses and TCP port numbers in the session table entries for HTTP traffic.

Session Limits

To avoid overloading a firewall, the ServerIron does not forward a packet to a firewall if either of the following conditions is true:

- The firewall already has the maximum allowed number of open sessions with the ServerIron. An open session is represented by a session entry. By default, a firewall can have up to one million session entries on the ServerIron. In a high-availability pair, the firewall can have up to one million combined on both ServerIrons. You can change the maximum number of sessions on an individual firewall basis to a number from 1 – 1,000,000.
- The firewall has already received the maximum allowed number of new sessions within the previous one-second interval. By default, the ServerIron will allow up to 65535 new sessions to the same firewall. The maximum includes TCP and UDP sessions combined. You can change the maximum number of sessions separately for TCP and UDP, to a value from 1 – 65535.

Session Aging

The ServerIron ages out inactive session entries. The aging mechanism differs depending on whether the session entry is a Layer 3 entry or a Layer 4 entry:

- Layer 3 session entries – The ServerIron uses the sticky age timer to age out Layer 3 session entries. The default sticky age is 5 minutes. You can change the sticky age to a value from 2 – 60 minutes.
  - To change the timer, enter the **server sticky-age** <num> command at the global CONFIG level of the CLI.
- Layer 4 session entries – The ServerIron clears a session entry that has TCP ports when the ServerIron receives a TCP FIN or RESET to end the session. For a TCP session that ends abnormally, the ServerIron uses the TCP age timer to age out the session. The ServerIron uses the UDP age timer to age out all UDP sessions. The default TCP age timer is 30 minutes. The default UDP age timer is 5 minutes. You can configure either timer to a value from 2 – 60 minutes.
  - To change the TCP age timer, enter the **server tcp-age** <num> command at the global CONFIG level of the CLI.

- To change the UDP age timer, enter the **server udp-age** <num> command at the global CONFIG level of the CLI.

NOTE: SLB uses the same values for the sticky age, TCP age, and UDP age timers. If you change a timer, the change applies to both SLB and FWLB.

Health Checks

The ServerIron regularly checks the health of the firewall and router paths, and of the applications on the firewalls, if you add applications to the firewall configurations.

Path Health Checks

One of the required FWLB parameters is a separate path from the ServerIron through each firewall to each of the ServerIrons on the other side of the firewall. A path to the ServerIron's gateway router also is required.

By default, the ServerIron performs a Layer 3 health check of each firewall and router path by sending an ICMP ping packet on each path.

- If the ServerIron receives a reply within the allowed amount of time, the ServerIron concludes that the path is good.
- If the ServerIron does not receive a reply within the allowed amount of time, the ServerIron concludes that the path is down.

By default, the ServerIron waits 400 milliseconds for a reply to an ICMP health check packet. If the reply does not arrive, the ServerIron makes two more attempts by default. Therefore, the total amount of time the ServerIron waits for a response is 1.2 seconds by default.

You can increase the total amount of time the ServerIron will wait for a response by increasing the number of attempts. The valid number of attempts is a value from 3 – 31.

The default maximum number of health check attempts is 3 and can be configured to a value from 3 – 31.

NOTE: You must configure the same path health check parameters on all the ServerIrons in the FWLB configuration. Otherwise, the paths will not pass the health checks.

Application Health Checks

When you add firewall configuration information to the ServerIron, you also can add information for individual application ports. Adding the application information is optional.

You can specify the following:

- The application's protocol (TCP or UDP) and port number
- The Layer 4 health check state (enabled or disabled) for the application

Adding an application port provides the following benefits:

- The ServerIron includes the source and destination port numbers for the application when it creates session entry. Thus, adding the application port provides more granular load balancing.
- The ServerIron checks the health of the TCP or UDP service used by the application, by sending a Layer 4 TCP or UDP health check to the firewall.

Layer 4 health checks are enabled by default. However, you can disable the Layer 4 health checks globally or on individual application on individual firewalls.

The ServerIron performs the Layer 4 TCP and UDP health checks as follows:

- TCP health check – The ServerIron checks the TCP port's health based on a TCP three-way handshake:

    - The ServerIron sends a TCP SYN packet to the port on the firewall.
    - The ServerIron expects the firewall to respond with a SYN ACK.
    - If the ServerIron receives the SYN ACK, the ServerIron sends a TCP RESET, satisfied that the TCP port is alive.
- UDP health check – The ServerIron sends a UDP packet with garbage (meaningless) data to the UDP port:
    - If the firewall responds with an ICMP “Port Unreachable” message, the ServerIron concludes that the port is not alive.
    - If the server does not respond at all, the ServerIron assumes that the port is alive and received the garbage data. Since UDP is a connectionless protocol, the ServerIron and other clients do not expect replies to data sent to a UDP port. Thus, lack of a response indicates a healthy port.

Configuring HA Active-Active FWLB

This section contains the following sections:

Overview of Active-Active FWLB

Active-Active operation provides redundancy in case a ServerIron becomes unavailable, while enhancing performance by using both ServerIrons to process and forward traffic.

NOTE: Active-Active operation is not the same thing as the always-active feature. The always-active feature is used to simplify the topology of high-availability FWLB configurations, and can be used in an Active-Active configuration.

Figure 5.1 shows an example of ServerIron Chassis device configured for high-availability FWLB.

Figure 5.1 HA FWLB for Layer 3 Firewalls

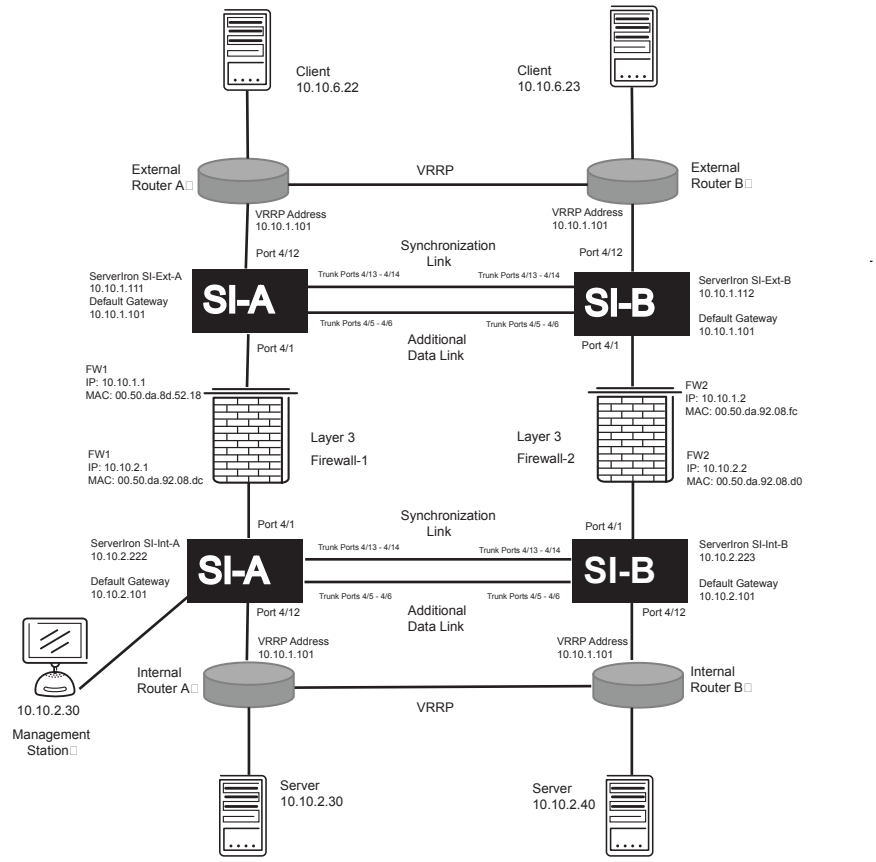

In this example, clients access the application servers on the private network through one of two routers, each of which is connected to a ServerIron. The ServerIrons create session entries for new traffic flows, including assignment of a firewall. The ServerIrons then use the session entries to forward subsequent traffic in the flow to the same firewall.

The ServerIrons on the private side of the network are connected to the application servers through routers. These ServerIrons also create session entries and use those entries for forwarding traffic to the servers and the server replies back to the clients.

Each pair of ServerIrons is connected by two trunk groups. One of the trunk groups is the synchronization link, and is used by the ServerIron to exchange session information, so that each ServerIron has a complete list of the sessions. If one of the ServerIrons becomes unavailable, the other ServerIron can continue FWLB service without interruption, even for existing sessions.

The other trunk group is an additional data link and allows for a simplified topology by eliminating the need for separate Layer 2 Switches between the ServerIrons and firewalls.

These links are not required to be trunk groups, but configuring them as trunk groups adds link-level redundancy to the overall redundant design.

The pairs of routers are configured with Virtual Router Redundancy Protocol (VRRP) to share the default gateway address used by the ServerIrons attached to the routers.

A management station attached to one of the ServerIrons on the private side of the firewalls provides Telnet management access to all four ServerIrons.

To implement the Active-Active FWLB configuration shown in Figure 5.1, perform the following tasks on each ServerIron.

Table 5.1: Configuration tasks – Active-Active FWLB

Task	See page...
Configure Global Parameters	
Configure the management IP address and default gateway	5-6
Configure the partner port, for the synchronization link	5-7
Configure the additional data link (the always-active link)	5-7
Configure the router port	5-7
Configure Firewall Parameters	
Define the firewalls and add them to the firewall group. When you define each firewall, optionally specify: • The TCP or UDP application ports on the firewall • The health check state (enabled by default) • The maximum total number of sessions • The maximum new session rate	5-8
Configure Firewall Group Parameters	
Change the load balancing method from least connections to least connections per application (optional)	5-10
Enable the active-active mode	5-11
Configure the paths and add static MAC entries for the firewall interfaces with the ServerIron	5-11
Configure the ServerIron to drop traffic when the firewall has reached its maximum number of sessions or maximum new session rate (optional)	5-12
Configure the ServerIron to forward a TCP data packet only if the ServerIron has already received a TCP SYN for the packet's source and destination addresses (optional)	5-12
Enable FWLB	
Assign all Web Switching Management Modules to a single WSM CPU for FWLB **Note**: This step is applicable only if you are running a software release earlier than 07.2.20 and the chassis is using more than one forwarding module.	5-12
Globally enable FWLB	5-13

Configuring the Management IP Address and Default Gateway

You must add a management IP address to the ServerIron and the IP address must be in the same sub-net as the ServerIron's interfaces with the Layer 3 firewalls.

For the default gateway address, specify the IP address on the router's interface with the ServerIron.


```
ServerIron(config)# ip address 10.10.1.111 255.255.255.0
ServerIron(config)# ip default-gateway 10.10.1.101
```

Syntax: ip address <ip-addr> <ip-mask>

or

Syntax: ip address <ip-addr>/<mask-bits>

Syntax: ip default-gateway <ip-addr>

Configuring the Partner Port

When you configure the ServerIron for IronClad FWLB, you need to specify the port number of the dedicated synchronization link between the ServerIron and its active-active partner. To specify the port, enter a command such as the following at the global CLI level:

```
ServerIron(config)# server fw-port 4/13
```

Syntax: [no] server fw-port <portnum>

If the link between the two ServerIrons is a trunk group (recommended for added redundancy), specify the port number of the primary port. The primary port is the first port in the trunk group.

Configuring the Additional Data Link (the Always-Active Link)

The default port-based VLAN, VLAN 1, contains all the ServerIron ports by default. In configurations such as the one shown in Figure 5.1 on page 5-5, the ports of the additional data link between the ServerIrons in each pair also are in this VLAN. For this type of configuration, you must perform the following configuration steps for the default VLAN:

- Disable the Spanning Tree Protocol (STP)
- Enable the always-active option

To disable STP and enable the always-active feature, enter the following commands:

```
ServerIron(config)# vlan 1
ServerIron(config-vlan-1)# no spanning-tree
ServerIron(config-vlan-1)# always-active
ServerIron(config-vlan-1)# exit
ServerIron(config)#
```

Syntax: [no] vlan <num>

Syntax: [no] spanning-tree

Syntax: [no] always-active

NOTE: To use the always-active feature, you also must enable the L2-fwall feature at the firewall group configuration level.

Configuring the Router Port

High-availability FWLB configurations require that you identify the ports on the ServerIron that are attached to the router(s).

To identify the router port, enter a command such as the following:

```
ServerIron(config)# server router-ports 4/12
```

Syntax: [no] server router-ports <portnum>

NOTE: To define multiple router ports on a switch, enter the port numbers, separated by blanks. You can enter up to eight router ports in a single command line. To enter more than eight ports, enter the **server router-ports** command again with the additional ports.

If the link is a trunk group, specify the port number of the primary port. The primary port is the first port in the trunk group.

Configuring the Firewalls

To configure a firewall, enter a name for the firewall and the IP address of its interface with the ServerIron. Optionally, you also can enter the following information:

- The TCP or UDP application ports on the firewall
- The health check state (enabled by default)
- The maximum total number of sessions
- The maximum new session rate

Adding the Firewalls

To configure the firewalls on ServerIron SI-Ext-A in Figure 5.1, enter commands such as the following:

```
ServerIron(config)# server fw-name FW1 10.10.10.1
ServerIron(config-rs-FW1)# port http
ServerIron(config-rs-FW1)# exit
ServerIron(config)# fw-name FW2 10.10.10.2
ServerIron(config-rs-FW2)# port http
ServerIron(config-rs-FW2)# exit
ServerIron(config)# server fw-group 2
ServerIron(config-tc-2)# fw-name FW1
ServerIron(config-tc-2)# fw-name FW2
```

Syntax: [no] server fw-name <string> <ip-addr>

This command adds a firewall.

Syntax: [no] port <tcp/udp-port> [no-health-check]

The <tcp/udp-port> parameter can be a number from 1 – 65535 or one of the following well-known port names:

- **dns** – port 53
- **ftp** – port 21. (Ports 20 and 21 both are FTP ports but in the ServerIron, the name “ftp” corresponds to port 21.)
- **http** – port 80
- **imap4** – port 143
- **ldap** – port 389
- **nntp** – port 119
- **ntp** – port 123
- **pop2** – port 109
- **pop3** – port 110
- **radius** – UDP port 1812
- **radius-old** – the ServerIron name for UDP port 1645, which is used in some older RADIUS implementations instead of port 1812
- **smtp** – port 25

- **snmp** – port 161
- **ssl** – port 443
- **telnet** – port 23
- **tftp** – port 69

The **no-health-check** parameter disables the Layer 4 path heath check for this application port. Layer 4 health checks are enabled by default.

Changing the Maximum Number of Sessions

To change the maximum number of sessions the firewall can have on the high-availability pair of ServerIrons, enter a command such as the following:

```
ServerIron(config-rs-FW1)# max-conn 145000
```

Syntax: [no] max-conn <num>

The <num> parameter specifies the maximum and can be from 1 – 1000000. This maximum applies to both the ServerIron and its high-availability partner.

NOTE: Most FWLB parameters, including this one, must be set to the same value on both ServerIrons in the high-availability pair.

NOTE: If you use the **max-conn** command for a firewall, the command specifies the maximum permissible number of connections that can be initiated from this ServerIron's direction on the firewall paths. The **max-conn** command does not limit the total number of connections that can exist on the ServerIron, which includes connections that come from the ServerIrons at the other ends of the firewall paths. For FWLB, the command to restrict the total number of connections that can exist on the ServerIron is **fw-exceed-max-drop**. See “Dropping Packets When a Firewall Reaches Its Limit” on page 5-12.

Connection Rate Control

Connection Rate Control (CRC) enables you to change the maximum number of new TCP sessions with the ServerIrons the firewall can have per second, enter a command such as the following:

```
ServerIron(config-rs-FW1)# max-tcp-conn-rate 1000
```

Syntax: [no] max-tcp-conn-rate <num>

Syntax: [no] max-udp-conn-rate <num>

The <num> parameter specifies the maximum number of connections per second and can be a number from 1 – 65535. The default is 65535.

Limiting the Number of New Connections for an Application

The following commands limit the rate of new connections per second to TCP port 80 on firewall FW1.

```
ServerIron(config)# server fw-name FW1 1.2.3.4
ServerIron(config-rs-FW1)# port http
ServerIron(config-rs-FW1)# port http max-tcp-conn-rate 800
```

Syntax: port <TCP/UDP-portnum> max-tcp-conn-rate <num>

Syntax: port <TCP/UDP-portnum> max-udp-conn-rate <num>

The **port** <TCP/UDP-portnum> parameter specifies the application port.

The <num> parameter specifies the maximum number of connections per second.

Adding the Firewalls to the Firewall Group

To add the firewalls to the firewall group, enter commands such as the following:

```
ServerIron(config-rs-FW1)# exit
ServerIron(config)# server fw-group-2
ServerIron(config-tc-2)# fw-name FW1
ServerIron(config-tc-2)# fw-name FW2
```

Syntax: server fw-group 2

This command changes the CLI to firewall group configuration level. The firewall group number is 2. Only one firewall group is supported.

Syntax: [no] fw-name <string>

This command adds a configured firewall to the firewall group.

Changing the Load-Balancing Method

By default, the ServerIron load balances firewall traffic flows by selecting the firewall with the lowest number of total connections. You can configure the ServerIron to load balance based on the lowest number of connections for the traffic flow's application.

For example, suppose a configuration has two firewalls (FW1 and FW2), and each firewall has two application ports defined (HTTP and SMTP). Also assume the following:

- FW1 has 10 HTTP connections and 80 SMTP connections.
- FW2 has 60 HTTP connections and 10 SMTP connections.

Using the default load balancing method, traffic for a new flow is load balanced to FW2, since this firewall has fewer total connections. This is true regardless of the application in the traffic. However, using the load balancing by application method, a new traffic flow carrying HTTP traffic is load balanced to FW1 instead of FW2, because FW1 has fewer HTTP connections. A new traffic flow for SMTP is load balanced to FW2, since FW2 has fewer SMTP connections.

To enable load balancing by application, enter the following command at the firewall group configuration level:

```
ServerIron(config-tc-2)# fw-predictor per-service-least-conn
```

Syntax: [no] fw-predictor total-least-conn | per-service-least-conn

The **total-least-conn** parameter load balances traffic based on the total number of connections only. This is the default.

The **per-service-least-conn** parameter load balances traffic based on the total number of connections for the traffic's application. This is valid for TCP or UDP applications.

Hashing Load Balance Metric in FWLB

NOTE: This feature applies to Releases 09.3.01 and later.

Fire Wall Load Balancing (FWLB) balances firewall traffic flows across multiple firewalls. Older ServerIron XL systems have always load balanced traffic to firewalls by hashing source IP and destination IP addresses. Optionally, if the **`hash-ports`** command was configured on the device, the hashing would include TCP source port and TCP destination port if the source or destination port was one of the ports listed with the **`hash-ports`** command. On ServerIron XL, hashing is the default load balancing scheme, and there are no sessions created when load balancing is performed this way.

Beginning with Release 09.3.01, hashing is a new metric added to ServerIron chassis devices' support of load balancing. For this feature, configure the **`fw-predictor hash`** command under the **`fw-group`**. When this command is configured, firewall selection is based on hashing of IP addresses (and optionally ports). However, unlike the ServerIron XL devices, chassis devices create sessions for the flow. The packet will be dropped if hashing picks a firewall and if either of the following is true:

- The **max-conn** reached for that firewall
- Connection rate is exceeded for the firewall or the firewall port

Connection rate can be specified at the FW level or a FW port level.

To configure the hashing features, enter commands such as the following:

```
SLB-SI-A(config)# server fw-group 2
SLB-SI-A(config-tc-2)# fw-predictor hash
```

Syntax: fw-predictor hash

Enabling the Active-Active Mode

To enable the active-active mode, enter a command such as the following at the firewall group configuration level:

```
ServerIron(config-tc-2)# sym-priority 1
```

Syntax: [no] sym-priority <num>

The **sym-priority** command enables the active-active mode. Since this command is also used for Symmetric SLB (SSLB), the command requires a number from 1 – 255. In SSLB, the number specifies the priority of the ServerIron and is used to determine the active ServerIron in the configuration. In active-active FWLB, both ServerIrons are active, so the number you enter does not affect the configuration. The CLI requires that you enter a number but the number is not used by the active-active FWLB configuration.

Configuring the Paths and Static MAC Address Entries

The paths go from one ServerIron to the other ServerIrons on the other side of each firewall. A path also goes to the router.

A path consists of the following parameters:

- The path ID – A number that identifies the path. The paths go from one ServerIron to the other through the firewalls. A path also goes to the router. On each ServerIron, the sequence of path IDs must be contiguous (with no gaps), starting with path ID 1. For example, path sequence 1, 2, 3, 4, 5 is valid. Path sequence 1, 3, 5 or 5, 4, 3, 2, 1 is not valid.
- The ServerIron port – The number of the port that connects the ServerIron to the firewall. If your configuration does not require static MAC entries, you can specify a dynamic port (65535) instead of the physical port number for firewall paths. Specifying the dynamic port allows the ServerIron to select the physical port for the path so you don't need to. You cannot specify the dynamic port for router paths. Router paths require the physical port number.
- The other ServerIron's IP address – The management address of the ServerIron on the other side of the firewall.
- The next-hop IP address – The IP address of the firewall interface connected to this ServerIron.

NOTE: FWLB paths must be fully meshed. When you configure a FWLB path on a ServerIron, make sure you also configure a reciprocal path on the ServerIron attached to the other end of the firewalls. For example, if you configure four paths to four separate firewalls, make sure you configure four paths on the other ServerIron.

NOTE: In addition to configuring the paths, some configurations require a static MAC entry for each firewall interface attached to the ServerIron. Each configuration example in this guide indicates whether the configuration requires static MAC entries. The static MAC entries are not required if the routers are using OSPF.

To configure paths for ServerIron SI-Ext-A in Figure 5.1 on page 5-5, enter the following commands:

```
ServerIron(config-tc-2)# fwall-info 1 4/1 10.10.2.222 10.10.1.1
ServerIron(config-tc-2)# fwall-info 2 4/5 10.10.2.222 10.10.1.2
ServerIron(config-tc-2)# fwall-info 3 4/1 10.10.2.223 10.10.1.1
ServerIron(config-tc-2)# fwall-info 4 4/5 10.10.2.223 10.10.1.2
```



```
ServerIron(config-tc-2)# fwall-info 5 4/12 10.10.1.101 10.10.1.101
```

Syntax: [no] fwall-info <path-num> <portnum> <other-ServerIron-ip> <next-hop-ip>

To configure the static MAC address entries for ServerIron SI-Ext-A in Figure 5.1, enter the following commands:

```
ServerIron(config-tc-2)# vlan 1
ServerIron(config-vlan-1)# static-mac-address 0050.da92.08fc ethernet 4/5 priority 1
router-type
ServerIron(config-vlan-1)# static-mac-address 0050.da8d.5218 ethernet 4/1 priority 1
router-type
```

Syntax: [no] static-mac-address <mac-addr> ethernet <portnum> [priority <0-7>] [host-type | router-type]

The priority can be 0 – 7 (0 is lowest and 7 is highest) for chassis devices and either normal-priority or high-priority for stackable devices. Use a priority higher than 0.

Use **router-type** for the entry type.

If you are using the always-active feature (by entering the always-active command in VLAN 1 for simplified Layer 2 topology), you also must enable the L2-Fwall feature by entering the following command:

```
ServerIron(config-tc-2)# l2-fwall
```

Syntax: [no] l2-fwall

Dropping Packets When a Firewall Reaches Its Limit

By default, if the ServerIron receives traffic that it needs to forward to a firewall, but the firewall already has the maximum number of sessions open or has exceeded its maximum connection rate, the ServerIron uses a hashing mechanism to select another firewall. The hashing mechanism selects another firewall based on the source and destination IP addresses and application port numbers in the packet.

If you want the ServerIron to drop the traffic instead of load balancing it using the hashing mechanism, enter a command such as the following:

```
ServerIron(config-tc-2)# fw-exceed-max-drop
```

Syntax: [no] fw-exceed-max-drop

The ServerIron drops traffic only until the firewall again has available sessions.

Restricting TCP Traffic to a Firewall to Established Sessions

By default, the ServerIron sends a properly addressed TCP data packet to a firewall regardless of whether the ServerIron has received a TCP SYN for the traffic flow. For example, if the ServerIron receives a TCP packet addressed to TCP port 8080 on IP address 1.1.1.1, the ServerIron forwards the packet to firewall connected to 1,1.1.1 regardless of whether the ServerIron has received a TCP SYN for the session between the packet's source and 1.1.1.1.

For tighter security, you can configure the ServerIron to forward a TCP data packet only if the ServerIron has already received a TCP SYN for the packet's traffic flow (source and destination addresses). For example, with the tighter security enabled, the ServerIron does not forward a TCP data packet to 1.1.1.1 unless the ServerIron has already received a TCP SYN for the session between the packet's source and 1.1.1.1.

To enable the tighter security, enter the following command at the global CONFIG level of the CLI:

```
ServerIron(config)# server fw-strict-sec
```

Syntax: [no] server fw-strict-sec

The feature applies globally to all TCP traffic received for FWLB.

Assigning FWLB Processing to a WSM CPU

By default, the software distributes processing for the forwarding modules in the chassis among the WSM CPUs. However, in software releases earlier than 07.2.20, all FWLB processing must be performed using the same WSM CPU.

NOTE: This step is applicable only if you are running a software release earlier than 07.2.20 and the chassis is using more than one forwarding module.

To display the WSM CPU allocations, enter the **show wsm-map** command.

To assign all forwarding modules to the same WSM CPU, enter commands such as the following:

```
ServerIron(config)# wsm wsm-map slot 3 wsm-slot 2 wsm-cpu 1
ServerIron(config)# wsm wsm-map slot 4 wsm-slot 2 wsm-cpu 1
```

These commands remap processing for the modules in slots 3 and 4 to WSM CPU 1 on the Web Switching Management Module in slot 2.

Syntax: wsm wsm-map <from-slotnum> wsm-slot <to-slotnum> wsm-cpu <cpunum>

The <from-slotnum> parameter specifies the slot that contains the forwarding module.

The <to-slotnum> parameter specifies the slot that contains the Web Switching Management Module.

The <cpunum> parameter specifies the WSM CPU on <to-slotnum> that will perform the processing. The WSM CPUs are numbered from 1 – 3.

Enabling FWLB

Enter the following commands at the global CONFIG level to enable FWLB for all TCP and UDP traffic:

```
ServerIron(config)# ip policy 1 fw tcp 0 global
ServerIron(config)# ip policy 2 fw udp 0 global
ServerIron(config)# write mem
```

Syntax: [no] ip policy <policy-num> fw tcp | udp 0 global

The <policy-num> value identifies the policy and can be a number from 1 – 64.

Each policy affects TCP or UDP traffic, so you must specify **tcp** or **udp**.

The value 0 following the **tcp | udp** parameter specifies that the policy applies to all ports of the specified type (TCP or UDP). In this command, “0” is equivalent to “any port number”. For FWLB, you must specify “0”.

Complete CLI Example

The following sections show the CLI commands for configuring the ServerIrons in Figure 5.1.

NOTE: The **wsm wsm-map** commands in these examples are required only if the chassis is using more than one forwarding module. Otherwise, do not use this command.

Commands on ServerIron SI-Ext-A

```
ServerIron> enable
ServerIron# configure terminal
ServerIron(config)# hostname SI-Ext-A
SI-Ext-A(config)# ip address 10.10.1.111 255.255.255.0
SI-Ext-A(config)# ip default-gateway 10.10.1.101
```

The commands above add a management IP address and default gateway address to the ServerIron. The IP address must be in the same sub-net as the ServerIron’s interfaces with the Layer 3 firewalls.

```
SI-Ext-A(config)# trunk switch ethernet 4/5 to 4/6
SI-Ext-A(config)# trunk switch ethernet 4/13 to 4/14
```

The commands above configure trunk groups for the synchronization link and the additional data link between this ServerIron and its high-availability partner.

```
SI-Ext-A(config)# vlan 1
SI-Ext-A(config-vlan-1)# always-active
```



```
SI-Ext-A(config-vlan-1)# no spanning-tree
SI-Ext-A(config-vlan-1)# exit
```

The commands above enable the always-active feature and disable the Spanning Tree Protocol (STP) in VLAN 1, which contains the ports that will carry the FWLB traffic.

```
SI-Ext-A(config)# vlan 2 name sync_link by port
SI-Ext-A(config-vlan-2)# untagged ethernet 4/13 to 4/14
SI-Ext-A(config-vlan-2)# no spanning-tree
SI-Ext-A(config-vlan-2)# exit
```

The commands above configure the ports for the synchronization link to the other ServerIron in a separate port-based VLAN. The separate VLAN is required. Add the ports as untagged ports.

```
SI-Ext-A(config)# server fw-port 4/13
```

The **server fw-port** command identifies the port that connects this ServerIron to its high-availability partner. If you use a trunk group, specify the first port in the group (the group's primary port).

```
SI-Ext-A(config)# server router-port 4/12
```

The **server router-port** command identifies the port that connects this ServerIron to its default gateway router.

```
SI-Ext-A(config)# server fw-name FW1 10.10.1.1
SI-Ext-A(config-rs-FW1)# port http
SI-Ext-A(config-rs-FW1)# exit
SI-Ext-A(config)# server fw-name FW2 10.10.1.2
SI-Ext-A(config-rs-FW2)# port http
SI-Ext-A(config-rs-FW2)# server fw-group 2
SI-Ext-A(config-tc-2)# fw-name FW1
SI-Ext-A(config-tc-2)# fw-name FW2
```

The commands above configure the firewalls and add them to the firewall group. Since an application port is configured on each firewall, the ServerIron will use Layer 4 sessions to load balance the firewall traffic for that application. The ServerIron will use Layer 3 sessions to load balance traffic for other applications.

```
SI-Ext-A(config-tc-2)# sym-priority 1
```

The command above enables the active-active mode. The number with the command is required by the CLI but is not used by FWLB. The CLI requires a number from 1 – 255 because the same command also is used to configure Symmetric SLB (SSLB), where the number determines the ServerIron's priority in the configuration.

```
SI-Ext-A(config-tc-2)# fwall-info 1 4/1 10.10.2.222 10.10.1.1
SI-Ext-A(config-tc-2)# fwall-info 2 4/5 10.10.2.222 10.10.1.2
SI-Ext-A(config-tc-2)# fwall-info 3 4/1 10.10.2.223 10.10.1.1
SI-Ext-A(config-tc-2)# fwall-info 4 4/5 10.10.2.223 10.10.1.2
SI-Ext-A(config-tc-2)# fwall-info 5 4/12 10.10.1.101 10.10.1.101
SI-Ext-A(config-tc-2)# l2-fwall
SI-Ext-A(config-tc-2)# exit
```

The commands above configure the data paths through the firewalls and to the default gateway router. The **l2-fwall** command is part of the always-active feature and is required if you use the **always-active** command.

```
SI-Ext-A(config)# vlan 1
SI-Ext-A(config-vlan-1)# static-mac-address 0050.da8d.5218 ethernet 4/1 priority 1
router-type
SI-Ext-A(config-vlan-1)# static-mac-address 0050.da92.08fc ethernet 4/5 priority 1
router-type
SI-Ext-A(config-vlan-1)# exit
```

The commands above add static entries to the ServerIron's MAC table for the firewall interfaces. Specify a priority higher than 0. You can specify a priority up to 7. The **router-type** parameter is required for FWLB.

```
SI-Ext-A(config)# wsm wsm-map slot 3 wsm-slot 2 wsm-cpu 1
SI-Ext-A(config)# wsm wsm-map slot 4 wsm-slot 2 wsm-cpu 1
```

The commands above remap the forwarding modules in slots 3 and 4 to WSM CPU 1 on the Web Switching Management Module in slot 2.

NOTE: The **wsm wsm-map** command is required only if the chassis is using more than one forwarding module.

```
SI-Ext-A(config)# ip policy 1 fw tcp 0 global
SI-Ext-A(config)# ip policy 2 fw udp 0 global
SI-Ext-A(config)# write memory
SI-Ext-A(config)# end
SI-Ext-A# reload
```

The commands above enable FWLB, save the configuration changes to the startup-config file, and reload the software.

NOTE: FWLB becomes active as soon as you enable it. However, you must reload the software to place the trunk group configuration into effect.

Commands on ServerIron SI-Ext-B

```
ServerIron> enable
ServerIron# configure terminal
ServerIron(config)# hostname SI-Ext-B
SI-Ext-B(config)# ip address 10.10.1.112 255.255.255.0
SI-Ext-B(config)# ip default-gateway 10.10.1.101
SI-Ext-B(config)# trunk switch ethernet 4/5 to 4/6
SI-Ext-B(config)# trunk switch ethernet 4/13 to 4/14
SI-Ext-B(config)# vlan 1
SI-Ext-B(config-vlan-1)# always-active
SI-Ext-B(config-vlan-1)# no spanning-tree
SI-Ext-B(config-vlan-1)# exit
SI-Ext-B(config)# vlan 2 name sync_link by port
SI-Ext-B(config-vlan-2)# untagged ethernet 4/13 to 4/14
SI-Ext-B(config-vlan-2)# no spanning-tree
SI-Ext-B(config-vlan-2)# exit
SI-Ext-B(config)# server fw-port 4/13
SI-Ext-B(config)# server router-ports 4/12
SI-Ext-B(config)# server fw-name FW1 10.10.1.1
SI-Ext-B(config-rs-FW1)# port http
SI-Ext-B(config-rs-FW1)# exit
SI-Ext-B(config)# server fw-name FW2 10.10.1.2
SI-Ext-B(config-rs-FW2)# port http
SI-Ext-B(config-rs-FW2)# server fw-group 2
SI-Ext-B(config-tc-2)# fw-name FW1
SI-Ext-B(config-tc-2)# fw-name FW2
SI-Ext-B(config-tc-2)# sym-priority 1
SI-Ext-B(config-tc-2)# fwall-info 1 4/5 10.10.2.222 10.10.1.1
SI-Ext-B(config-tc-2)# fwall-info 2 4/1 10.10.2.222 10.10.1.2
SI-Ext-B(config-tc-2)# fwall-info 3 4/5 10.10.2.223 10.10.1.1
SI-Ext-B(config-tc-2)# fwall-info 4 4/1 10.10.2.223 10.10.1.2
SI-Ext-B(config-tc-2)# fwall-info 5 4/12 10.10.1.101 10.10.1.101
SI-Ext-B(config-tc-2)# l2-fwall
SI-Ext-B(config-tc-2)# exit
SI-Ext-B(config)# vlan 1
SI-Ext-B(config-vlan-1)# static-mac-address 0050.da8d.5218 ethernet 4/5 priority 1
router-type
SI-Ext-B(config-vlan-1)# static-mac-address 0050.da92.08fc ethernet 4/1 priority 1
router-type
SI-Ext-B(config-vlan-1)# exit
SI-Ext-B(config)# wsm wsm-map slot 3 wsm-slot 2 wsm-cpu 1
SI-Ext-B(config)# wsm wsm-map slot 4 wsm-slot 2 wsm-cpu 1
SI-Ext-B(config)# ip policy 1 fw tcp 0 global
SI-Ext-B(config)# ip policy 2 fw udp 0 global
```



```
SI-Ext-B(config)# write memory
SI-Ext-B(config)# end
SI-Ext-B# reload
```

Commands on ServerIron SI-Int-A

```
ServerIron> enable
ServerIron# configure terminal
ServerIron(config)# hostname SI-Int-A
SI-Int-A(config)# ip address 10.10.2.222 255.255.255.0
SI-Int-A(config)# ip default-gateway 10.10.2.101
SI-Int-A(config)# trunk switch ethernet 4/5 to 4/6
SI-Int-A(config)# trunk switch ethernet 4/13 to 4/14
SI-Int-A(config)# vlan 1
SI-Int-A(config-vlan-1)# always-active
SI-Int-A(config-vlan-1)# no spanning-tree
SI-Int-A(config-vlan-1)# exit
SI-Int-A(config)# vlan 2 name sync_link by port
SI-Int-A(config-vlan-2)# untagged ethernet 4/13 to 4/14
SI-Int-A(config-vlan-2)# no spanning-tree
SI-Int-A(config-vlan-2)# exit
SI-Int-A(config)# server fw-port 4/13
SI-Int-A(config)# server router-ports 4/12
SI-Int-A(config)# server fw-name FW1 10.10.2.1
SI-Int-A(config-rs-FW1)# port http
SI-Int-A(config-rs-FW1)# exit
SI-Int-A(config)# server fw-name FW2 10.10.2.2
SI-Int-A(config-rs-FW2)# port http
SI-Int-A(config-rs-FW2)# server fw-group 2
SI-Int-A(config-tc-2)# fw-name FW1
SI-Int-A(config-tc-2)# fw-name FW2
SI-Int-A(config-tc-2)# sym-priority 1
SI-Int-A(config-tc-2)# fwall-info 1 4/1 10.10.1.111 10.10.2.1
SI-Int-A(config-tc-2)# fwall-info 2 4/5 10.10.1.111 10.10.2.2
SI-Int-A(config-tc-2)# fwall-info 3 4/1 10.10.1.112 10.10.2.1
SI-Int-A(config-tc-2)# fwall-info 4 4/5 10.10.1.112 10.10.2.2
SI-Int-A(config-tc-2)# fwall-info 5 4/12 10.10.2.101 10.10.2.101
SI-Int-A(config-tc-2)# l2-fwall
SI-Int-A(config-tc-2)# exit
SI-Int-A(config)# vlan 1
SI-Int-A(config-vlan-1)# static-mac-address 0050.da92.08dc ethernet 4/1 priority 1
router-type
SI-Int-A(config-vlan-1)# static-mac-address 0050.da92.08d0 ethernet 4/5 priority 1
router-type
SI-Int-A(config-vlan-1)# exit
SI-Int-A(config)# wsm wsm-map slot 3 wsm-slot 2 wsm-cpu 1
SI-Int-A(config)# wsm wsm-map slot 4 wsm-slot 2 wsm-cpu 1
SI-Int-A(config)# ip policy 1 fw tcp 0 global
SI-Int-A(config)# ip policy 2 fw udp 0 global
SI-Int-A(config)# write memory
SI-Int-A(config)# end
SI-Int-A# reload
```

Commands on ServerIron SI-Int-B

```
ServerIron> enable
ServerIron# configure terminal
ServerIron(config)# hostname SI-Int-B
SI-Int-B(config)# ip address 10.10.2.223 255.255.255.0
SI-Int-B(config)# ip default-gateway 10.10.2.101
```



```
SI-Int-B(config)# trunk switch ethernet 4/5 to 4/6
SI-Int-B(config)# trunk switch ethernet 4/13 to 4/14
SI-Int-B(config)# vlan 1
SI-Int-B(config-vlan-1)# always-active
SI-Int-B(config-vlan-1)# no spanning-tree
SI-Int-B(config-vlan-1)# exit
SI-Int-B(config)# vlan 2 name sync_link by port
SI-Int-B(config-vlan-2)# untagged ethernet 4/13 to 4/14
SI-Int-B(config-vlan-2)# no spanning-tree
SI-Int-B(config-vlan-2)# exit
SI-Int-B(config)# server fw-port 4/13
SI-Int-B(config)# server router-ports 4/12
SI-Int-B(config)# server fw-name FW1 10.10.2.1
SI-Int-B(config-rs-FW1)# port http
SI-Int-B(config-rs-FW1)# exit
SI-Int-B(config)# server fw-name FW2 10.10.2.2
SI-Int-B(config-rs-FW2)# port http
SI-Int-B(config-rs-FW2)# server fw-group 2
SI-Int-B(config-tc-2)# fw-name FW1
SI-Int-B(config-tc-2)# fw-name FW2
SI-Int-B(config-tc-2)# sym-priority 1
SI-Int-B(config-tc-2)# fwall-info 1 4/5 10.10.1.111 10.10.2.1
SI-Int-B(config-tc-2)# fwall-info 2 4/1 10.10.1.111 10.10.2.2
SI-Int-B(config-tc-2)# fwall-info 3 4/5 10.10.1.112 10.10.2.1
SI-Int-B(config-tc-2)# fwall-info 4 4/1 10.10.1.112 10.10.2.2
SI-Int-B(config-tc-2)# fwall-info 5 4/12 10.10.2.101 10.10.2.101
SI-Int-B(config-tc-2)# l2-fwall
SI-Int-B(config-tc-2)# exit
SI-Int-B(config)# vlan 1
SI-Int-B(config-vlan-1)# static-mac-address 0050.da92.08dc ethernet 4/5 priority 1
router-type
SI-Int-B(config-vlan-1)# static-mac-address 0050.da92.08d0 ethernet 4/1 priority 1
router-type
SI-Int-B(config-vlan-1)# exit
SI-Int-B(config)# wsm wsm-map slot 3 wsm-slot 2 wsm-cpu 1
SI-Int-B(config)# wsm wsm-map slot 4 wsm-slot 2 wsm-cpu 1
SI-Int-B(config)# ip policy 1 fw tcp 0 global
SI-Int-B(config)# ip policy 2 fw udp 0 global
SI-Int-B(config)# write memory
SI-Int-B(config)# end
SI-Int-B# reload
```

Configuring New Active-Active HA FWLB

NOTE: This new configuration applies to Releases 09.3.01 and later.

The following Active-Active FWLB configuration has been tested recently. The commands presented below are documented in the *ServerIron Chassis L4-7 Software Configuration Guide*, except for the **other ip** command. The **other ip** command interprets the synch messages if firewall IP addresses are different on different ServerIrons.

Example

The following configuration and diagram is example of how active-active FWLB is configured in release 09.3.01.

Figure 5.2 Active-Active FWLB Topology

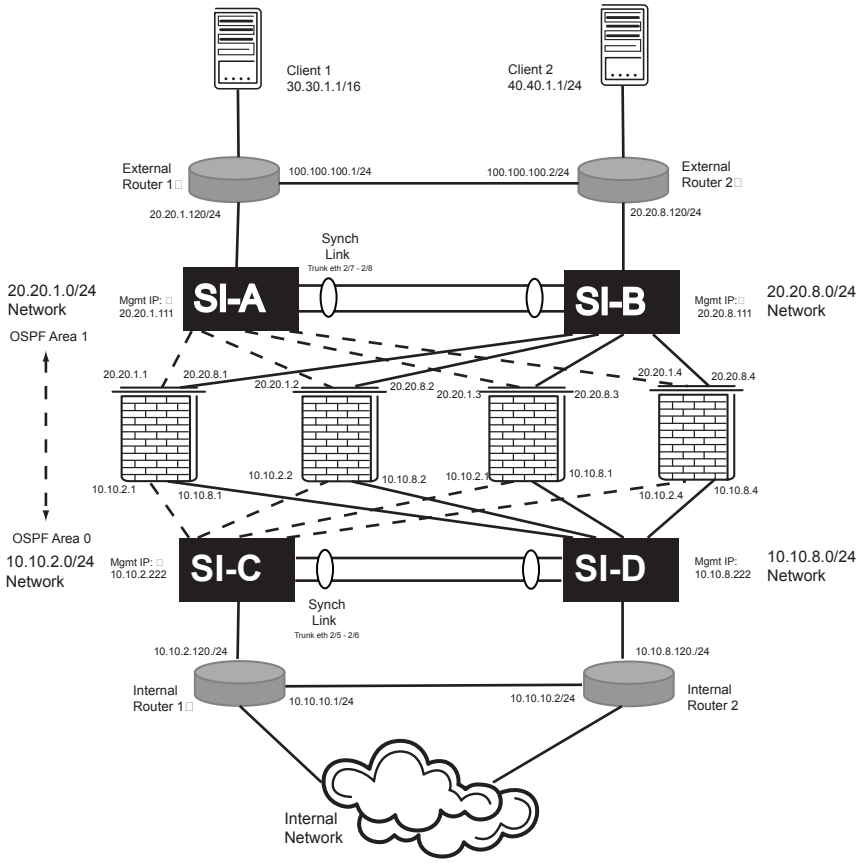

Notes about the configuration:

- The code was tested on WSM6 and JetCore modules.
- This topology looks similar to the ServerIronI-XL's active-standby topology, but FWLB ServerIrons work in active-active mode. Firewall paths will be up on the both the ServerIrons and both ServerIrons can do FWLB.
- The **always-active** command is configured under VLAN 1. This command should not be configured under synch ports vlan.
- In Chassis releases, stateful algorithm is used for FWLB; therefore, the ServerIron needs to synchronize sessions with its partner ServerIron to support stateful fail-over in high availability FWLB configurations.
- In the topology presented in this section, IP addresses of firewalls are different on each ServerIron. Use the **other-ip** command under firewall configuration level to identify the partner ServerIron's firewall address.
- This topology assumes that OSPF is running on firewalls, external routers, and internal routers. These devices exchange OSPF messages (multicast packets) among them. When a ServerIron is in state 3, it will block multicast packets. In the attached topology, if Ext-SI-B is in state 3, it will block the OSPF multicast packets sent by the firewalls and Ext-Router-2 to prevent Ext-Router-2 and the firewalls from learning OSPF routes through each other. Ext-Router-2 learns the OSPF routes of internal networks through Ext-Router-1. So all the external traffic will be going to Ext-SI-A.
- If the design requires ServerIron (in state 3) not to block multicast packets, the **server fw-allow-multicast** must be configured on the ServerIrons. When the command is configured, the external routers can learn the OSPF routes from the firewalls and traffic can go to both ServerIrons.

External ServerIron Standby A (Ext-SI-A) Configuration

```
SI-StandbyA(config)# module 1 bi-0-port-wsm2-management-module
```



```
SI-StandbyA(config)# module 2 bi-jc-8-port-gig-module
SI-StandbyA(config)# module 3 bi-jc-16-port-gig-copper-module
SI-StandbyA(config)# trunk switch ethernet 2/7 to 2/8
SI-StandbyA(config)# server fw-port 2/7
SI-StandbyA(config)# server router-ports ethernet 2/1
SI-StandbyA(config)# server fw-name fw1 20.20.1.1
SI-StandbyA(config-rs-FW1)# other-ip 20.20.8.1
SI-StandbyA(config-rs-FW1)# port http
SI-StandbyA(config-rs-FW1)# port http no-health-check
SI-StandbyA(config-rs-FW1)# port http url "HEAD /"
SI-StandbyA(config-rs-FW1)# exit
SI-StandbyA(config)# server fw-name fw2 20.20.1.2
SI-StandbyA(config-rs-FW2)# other-ip 20.20.8.2
SI-StandbyA(config-rs-FW2)# port http
SI-StandbyA(config-rs-FW2)# port http no-health-check
SI-StandbyA(config-rs-FW2)# port http url "HEAD /"
SI-StandbyA(config-rs-FW2)# exit
SI-StandbyA(config)# server fw-name fw3 20.20.1.3
SI-StandbyA(config-rs-FW3)# other-ip 20.20.8.3
SI-StandbyA(config-rs-FW3)# port http
SI-StandbyA(config-rs-FW3)# port http no-health-check
SI-StandbyA(config-rs-FW3)# port http url "HEAD /"
SI-StandbyA(config-rs-FW3)# exit
SI-StandbyA(config)# server fw-name fw4 20.20.1.4
SI-StandbyA(config-rs-FW4)# other-ip 20.20.8.4
SI-StandbyA(config-rs-FW4)# port http
SI-StandbyA(config-rs-FW4)# port http no-health-check
SI-StandbyA(config-rs-FW4)# port http url "HEAD /"
SI-StandbyA(config-rs-FW4)# server fw-group 2
SI-StandbyA(config-tc-2)# l2-fwall
SI-StandbyA(config-tc-2)# sym-priority 250
SI-StandbyA(config-tc-2)# fw-name fw1
SI-StandbyA(config-tc-2)# fw-name fw2
SI-StandbyA(config-tc-2)# fw-name fw3
SI-StandbyA(config-tc-2)# fw-name fw4
SI-StandbyA(config-tc-2)# fwall-info 1 3/1 10.10.2.222 20.20.1.1
SI-StandbyA(config-tc-2)# fwall-info 2 3/2 10.10.2.222 20.20.1.2
SI-StandbyA(config-tc-2)# fwall-info 3 3/3 10.10.2.222 20.20.1.3
SI-StandbyA(config-tc-2)# fwall-info 4 3/4 10.10.2.222 20.20.1.4
SI-StandbyA(config-tc-2)# fwall-info 5 3/1 10.10.8.222 20.20.1.1
SI-StandbyA(config-tc-2)# fwall-info 6 3/2 10.10.8.222 20.20.1.2
SI-StandbyA(config-tc-2)# fwall-info 7 3/3 10.10.8.222 20.20.1.3
SI-StandbyA(config-tc-2)# fwall-info 8 3/4 10.10.8.222 20.20.1.4
SI-StandbyA(config-tc-2)# fwall-info 9 2/1 20.20.1.120 20.20.1.120
SI-StandbyA(config-tc-2)# fw-predictor per-service-least-conn
SI-StandbyA(config-tc-2)# exit
SI-StandbyA(config)# vlan 1 name DEFAULT-VLAN by port
SI-StandbyA(config-vlan-1)# always-active
SI-StandbyA(config-vlan-1)# no spanning-tree
SI-StandbyA(config-vlan-1)# static-mac-address 0004.80ed.17b4 ethernet 3/1 priority
1 router-type
SI-StandbyA(config-vlan-1)# static-mac-address 0004.80f0.4b3c ethernet 3/2 priority
1 router-type
SI-StandbyA(config-vlan-1)# static-mac-address 0004.80ed.1368 ethernet 3/3 priority
1 router-type
SI-StandbyA(config-vlan-1)# static-mac-address 0004.80eb.5294 ethernet 3/4 priority
1 router-type
SI-StandbyA(config-vlan-1)# exit
SI-StandbyA(config)# vlan 999 by port
```



```
SI-StandbyA(config-vlan-999)# untagged ethernet 2/7 to 2/8
SI-StandbyA(config-vlan-999)# no spanning-tree
SI-StandbyA(config-vlan-999)# exit
SI-StandbyA(config)# hostname Ext-SI-A
SI-StandbyA(config)# ip address 20.20.1.111 255.255.255.0
SI-StandbyA(config)# ip default-gateway 20.20.1.120
SI-StandbyA(config)# auto-cam-repaint
SI-StandbyA(config)# pram-write-retry
SI-StandbyA(config)# write memory
SI-StandbyA(config)# end
SI-StandbyA(config)# end reload
```

External ServerIron Standby B (Ext-SI-B) Configuration

```
SI-StandbyB(config)# module 1 bi-0-port-wsm2-management-module
SI-StandbyB(config)# module 2 bi-jc-8-port-gig-module
SI-StandbyB(config)# module 3 bi-jc-16-port-gig-copper-module
SI-StandbyB(config)# trunk switch ethernet 2/7 to 2/8
SI-StandbyB(config)# server fw-port 2/7
SI-StandbyB(config)# server router-ports ethernet 2/1
SI-StandbyB(config)# server fw-name fw1 20.20.8.1
SI-StandbyB(config-rs-FW1)# other-ip 20.20.1.1
SI-StandbyB(config-rs-FW1)# port http
SI-StandbyB(config-rs-FW1)# port http no-health-check
SI-StandbyB(config-rs-FW1)# port http url "HEAD /"
SI-StandbyB(config-rs-FW1)# exit
SI-StandbyB(config)#server fw-name fw2 20.20.8.2
SI-StandbyB(config-rs-FW2)# other-ip 20.20.1.2
SI-StandbyB(config-rs-FW2)# port http
SI-StandbyB(config-rs-FW2)# port http no-health-check
SI-StandbyB(config-rs-FW2)# port http url "HEAD /"
SI-StandbyB(config-rs-FW2)# exit
SI-StandbyB(config)# server fw-name fw3 20.20.8.3
SI-StandbyB(config-rs-FW3)# other-ip 20.20.1.3
SI-StandbyB(config-rs-FW3)# port http
SI-StandbyB(config-rs-FW3)# port http no-health-check
SI-StandbyB(config-rs-FW3)# port http url "HEAD /"
SI-StandbyB(config-rs-FW3)# exit
SI-StandbyB(config)# server fw-name fw4 20.20.8.4
SI-StandbyB(config-rs-FW4)# other-ip 20.20.1.4
SI-StandbyB(config-rs-FW4)# port http
SI-StandbyB(config-rs-FW4)# port http no-health-check
SI-StandbyB(config-rs-FW4)# port http url "HEAD /"
SI-StandbyB(config-rs-FW4)# exit
SI-StandbyB(config-rs-FW4)# server fw-group 2
SI-StandbyB(config-rs-tc-2)# l2-fwall
SI-StandbyB(config-rs-tc-2)# sym-priority 10
SI-StandbyB(config-rs-tc-2)# fw-name fw1
SI-StandbyB(config-rs-tc-2)# fw-name fw2
SI-StandbyB(config-rs-tc-2)# fw-name fw3
SI-StandbyB(config-rs-tc-2)# fw-name fw4
SI-StandbyB(config-rs-tc-2)# fwall-info 1 3/1 10.10.8.222 20.20.8.1
SI-StandbyB(config-rs-tc-2)# fwall-info 2 3/2 10.10.8.222 20.20.8.2
SI-StandbyB(config-rs-tc-2)# fwall-info 3 3/3 10.10.8.222 20.20.8.3
SI-StandbyB(config-rs-tc-2)# fwall-info 4 3/4 10.10.8.222 20.20.8.4
SI-StandbyB(config-rs-tc-2)# fwall-info 5 3/1 10.10.2.222 20.20.8.1
SI-StandbyB(config-rs-tc-2)# fwall-info 6 3/2 10.10.2.222 20.20.8.2
SI-StandbyB(config-rs-tc-2)# fwall-info 7 3/3 10.10.2.222 20.20.8.3
SI-StandbyB(config-rs-tc-2)# fwall-info 8 3/4 10.10.2.222 20.20.8.4
```



```
SI-StandbyB(config-rs-tc-2)# fwall-info 9 2/1 20.20.8.120 20.20.8.120
SI-StandbyB(config-rs-tc-2)# fw-predictor per-service-least-conn
SI-StandbyB(config-rs-tc-2)# exit
SI-StandbyB(config)# vlan 1 name DEFAULT-VLAN by port
SI-StandbyB(config-vlan-1)# always-active
SI-StandbyB(config-vlan-1)# no spanning-tree
SI-StandbyB(config-vlan-1)# static-mac-address 0004.80ed.17b4 ethernet 3/1 priority
1 router-type
SI-StandbyB(config-vlan-1)# static-mac-address 0004.80f0.4b3c ethernet 3/2 priority
1 router-type
SI-StandbyB(config-vlan-1)# static-mac-address 0004.80ed.1368 ethernet 3/3 priority
1 router-type
SI-StandbyB(config-vlan-1)# static-mac-address 0004.80eb.5294 ethernet 3/4 priority
1 router-type
SI-StandbyB(config-vlan-1)# exit
SI-StandbyB(config)# vlan 999 by port
SI-StandbyB(config-vlan-999)# untagged ethe 2/7 to 2/8
SI-StandbyB(config-vlan-999)# no spanning-tree
SI-StandbyB(config-vlan-999)# exit
SI-StandbyB(config)# hostname Ext-SI-B
SI-StandbyB(config)# ip address 20.20.8.111 255.255.255.0
SI-StandbyB(config)# ip default-gateway 20.20.8.120
SI-StandbyB(config)# auto-cam-repaint
SI-StandbyB(config)# pram-write-retry
SI-StandbyB(config)# write memory
SI-StandbyB(config)# end
SI-StandbyB(config)# reload
```

Internal ServerIron C (Int-SI-C) Configuration

```
SI-ActiveC(config)# module 1 bi-0-port-wsm2-management-module
SI-ActiveC(config)# module 2 bi-jc-8-port-gig-module
SI-ActiveC(config)# module 3 bi-jc-16-port-gig-copper-module
SI-ActiveC(config)# trunk switch ethe 2/5 to 2/6
SI-ActiveC(config)# server fw-port 2/5
SI-ActiveC(config)# server router-ports ethernet 2/1
SI-ActiveC(config)# server fw-name fw1 10.10.2.1
SI-ActiveC(config-rs-FW1)# other-ip 10.10.8.1
SI-ActiveC(config-rs-FW1)# port http
SI-ActiveC(config-rs-FW1)# port http no-health-check
SI-ActiveC(config-rs-FW1)# port http url "HEAD /"
SI-ActiveC(config-rs-FW1)# exit
SI-ActiveC(config)# server fw-name fw2 10.10.2.2
SI-ActiveC(config-rs-FW2)# other-ip 10.10.8.2
SI-ActiveC(config-rs-FW2)# port http
SI-ActiveC(config-rs-FW2)# port http no-health-check
SI-ActiveC(config-rs-FW2)# port http url "HEAD /"
SI-ActiveC(config-rs-FW2)# exit
SI-ActiveC(config)# server fw-name fw3 10.10.2.3
SI-ActiveC(config-rs-FW3)# other-ip 10.10.8.3
SI-ActiveC(config-rs-FW3)# port http
SI-ActiveC(config-rs-FW3)# port http no-health-check
SI-ActiveC(config-rs-FW3)# port http url "HEAD /"
SI-ActiveC(config-rs-FW3)# exit
SI-ActiveC(config)# server fw-name fw4 10.10.2.4
SI-ActiveC(config-rs-FW4)# other-ip 10.10.8.4
SI-ActiveC(config-rs-FW4)# port http
SI-ActiveC(config-rs-FW4)# port http no-health-check
SI-ActiveC(config-rs-FW4)# port http url "HEAD /"
```



```
SI-ActiveC(config-rs-FW4)# exit
SI-ActiveC(config-rs-FW4)# server fw-group 2
SI-ActiveC(config-tc-2)# l2-fwall
SI-ActiveC(config-tc-2)# sym-priority 250
SI-ActiveC(config-tc-2)# fw-name fw1
SI-ActiveC(config-tc-2)# fw-name fw2
SI-ActiveC(config-tc-2)# fw-name fw3
SI-ActiveC(config-tc-2)# fw-name fw4
SI-ActiveC(config-tc-2)# fwall-info 1 3/1 20.20.1.111 10.10.2.1
SI-ActiveC(config-tc-2)# fwall-info 2 3/2 20.20.1.111 10.10.2.2
SI-ActiveC(config-tc-2)# fwall-info 3 3/3 20.20.1.111 10.10.2.3
SI-ActiveC(config-tc-2)# fwall-info 4 3/4 20.20.1.111 10.10.2.4
SI-ActiveC(config-tc-2)# fwall-info 5 3/1 20.20.8.111 10.10.2.1
SI-ActiveC(config-tc-2)# fwall-info 6 3/2 20.20.8.111 10.10.2.2
SI-ActiveC(config-tc-2)# fwall-info 7 3/3 20.20.8.111 10.10.2.3
SI-ActiveC(config-tc-2)# fwall-info 8 3/4 20.20.8.111 10.10.2.4
SI-ActiveC(config-tc-2)# fwall-info 9 2/1 10.10.2.120 10.10.2.120
SI-ActiveC(config-tc-2)# fw-predictor per-service-least-conn
SI-ActiveC(config-tc-2)# exit
SI-ActiveC(config)# vlan 1 name DEFAULT-VLAN by port
SI-ActiveC(config-vlan-1)# always-active
SI-ActiveC(config-vlan-1)# no spanning-tree
SI-ActiveC(config-vlan-1)# static-mac-address 0004.80ed.17b4 ethernet 3/1 priority 1
router-type
SI-ActiveC(config-vlan-1)# static-mac-address 0004.80f0.4b3c ethernet 3/2 priority 1
router-type
SI-ActiveC(config-vlan-1)# static-mac-address 0004.80ed.1368 ethernet 3/3 priority 1
router-type
SI-ActiveC(config-vlan-1)# static-mac-address 0004.80eb.5294 ethernet 3/4 priority 1
router-type
SI-ActiveC(config-vlan-1)# exit
SI-ActiveC(config)# vlan 999 by port
SI-ActiveC(config-vlan-999)# untagged ethe 2/5 to 2/8
SI-ActiveC(config-vlan-999)# no spanning-tree
SI-ActiveC(config-vlan-999)# exit
SI-ActiveC(config)# hostname Int-SI-C
SI-ActiveC(config)# ip address 10.10.2.222 255.255.255.0
SI-ActiveC(config)# ip default-gateway 10.10.2.120
SI-ActiveC(config)# auto-cam-repaint
SI-ActiveC(config)# pram-write-retry
SI-ActiveC(config)# write mem
SI-ActiveC(config)# reload
SI-ActiveC(config)# end
```

Internal ServerIron D (Int-SI-D) Configuration

```
SI-ActiveD(config)# module 1 bi-0-port-wsm2-management-module
SI-ActiveD(config)# module 2 bi-jc-8-port-gig-module
SI-ActiveD(config)# module 3 bi-jc-16-port-gig-copper-module
SI-ActiveD(config)# trunk switch ethe 2/5 to 2/6
SI-ActiveD(config)# server fw-port 2/5
SI-ActiveD(config)# server router-ports ethernet 2/1
SI-ActiveD(config)# server fw-name fw1 10.10.8.1
SI-ActiveD(config-rs-FW1)# other-ip 10.10.2.1
SI-ActiveD(config-rs-FW1)# port http
SI-ActiveD(config-rs-FW1)# port http no-health-check
SI-ActiveD(config-rs-FW1)# port http url "HEAD /"
SI-ActiveD(config-rs-FW1)# exit
SI-ActiveD(config)# server fw-name fw2 10.10.8.2
```



```
SI-ActiveD(config-rs-FW2)# other-ip 10.10.2.2
SI-ActiveD(config-rs-FW2)# port http
SI-ActiveD(config-rs-FW2)# port http no-health-check
SI-ActiveD(config-rs-FW2)# port http url "HEAD /"
SI-ActiveD(config-rs-FW2)# exit
SI-ActiveD(config)# server fw-name fw3 10.10.8.3
SI-ActiveD(config-rs-FW3)# other-ip 10.10.2.3
SI-ActiveD(config-rs-FW3)# port http
SI-ActiveD(config-rs-FW3)# port http no-health-check
SI-ActiveD(config-rs-FW3)# port http url "HEAD /"
SI-ActiveD(config-rs-FW3)#
SI-ActiveD(config)# server fw-name fw4 10.10.8.4
SI-ActiveD(config-rs-FW4)# other-ip 10.10.2.4
SI-ActiveD(config-rs-FW4)# port http
SI-ActiveD(config-rs-FW4)# port http no-health-check
SI-ActiveD(config-rs-FW4)# port http url "HEAD /"
SI-ActiveD(config-rs-FW4)# exit
SI-ActiveD(config-rs-FW4)# server fw-group 2
SI-ActiveD(config-tc-2)# l2-fwall
SI-ActiveD(config-tc-2)# sym-priority 10
SI-ActiveD(config-tc-2)# fw-name fw1
SI-ActiveD(config-tc-2)# fw-name fw2
SI-ActiveD(config-tc-2)# fw-name fw3
SI-ActiveD(config-tc-2)# fw-name fw4
SI-ActiveD(config-tc-2)# fwall-info 1 3/1 20.20.8.111 10.10.8.1
SI-ActiveD(config-tc-2)# fwall-info 2 3/2 20.20.8.111 10.10.8.2
SI-ActiveD(config-tc-2)# fwall-info 3 3/3 20.20.8.111 10.10.8.3
SI-ActiveD(config-tc-2)# fwall-info 4 3/4 20.20.8.111 10.10.8.4
SI-ActiveD(config-tc-2)# fwall-info 5 3/1 20.20.1.111 10.10.8.1
SI-ActiveD(config-tc-2)# fwall-info 6 3/2 20.20.1.111 10.10.8.2
SI-ActiveD(config-tc-2)# fwall-info 7 3/3 20.20.1.111 10.10.8.3
SI-ActiveD(config-tc-2)# fwall-info 8 3/4 20.20.1.111 10.10.8.4
SI-ActiveD(config-tc-2)# fwall-info 9 2/1 10.10.8.120 10.10.8.120
SI-ActiveD(config-tc-2)# fw-predictor per-service-least-conn
SI-ActiveD(config-tc-2)# exit
SI-ActiveD(config)# vlan 1 name DEFAULT-VLAN by port
SI-ActiveD(config-vlan-1)# always-active
SI-ActiveD(config-vlan-1)# no spanning-tree
SI-ActiveD(config-vlan-1)# static-mac-address 0004.80ed.17b4 ethernet 3/1 priority 1
router-type
SI-ActiveD(config-vlan-1)# static-mac-address 0004.80f0.4b3c ethernet 3/2 priority 1
router-type
SI-ActiveD(config-vlan-1)# static-mac-address 0004.80ed.1368 ethernet 3/3 priority 1
router-type
SI-ActiveD(config-vlan-1)# static-mac-address 0004.80eb.5294 ethernet 3/4 priority 1
router-type
SI-ActiveD(config-vlan-1)# exit
SI-ActiveD(config)# vlan 999 by port
SI-ActiveD(config-vlan-999)# untagged ethe 2/5 to 2/8
SI-ActiveD(config-vlan-999)# no spanning-tree
SI-ActiveD(config-vlan-999)# exit
SI-ActiveD(config)# hostname Int-SI-D
SI-ActiveD(config)# ip address 10.10.8.222 255.255.255.0
SI-ActiveD(config)# ip default-gateway 10.10.8.120
SI-ActiveD(config)# auto-cam-repaint
SI-ActiveD(config)# pram-write-retry
SI-ActiveD(config)# write memory
SI-ActiveD(config)# reload
SI-ActiveD(config)# end
```


Configuring Active-Active HA FWLB with VRRP

NOTE: Layer 3 routing is supported only on ServerIron Chassis devices running software release 08.0.00 or later.

This section shows examples of commonly used ServerIron IronClad FWLB deployments with Layer 3 configurations. The ServerIrons in these examples perform Layer 3 routing in addition to Layer 2 and Layer 4 – 7 switching.

Generally, the steps for configuring Layer 4 – 7 features on a ServerIron running Layer 3 are similar to the steps on a ServerIron that is not running Layer 3. The examples focus on the Layer 3 aspects of the configurations.

This section contains the following configuration examples:

- "Overview of Active-Active FWLB with VRRP" on page 5-24

NOTE: The IronClad FWLB configurations shown in these examples are the ones that are supported. If you need to use the ServerIron's Layer 3 routing support in a FWLB configuration that is not shown, contact Brocade Communications Systems.

Overview of Active-Active FWLB with VRRP

Figure 5.3 shows an example of an active-active FWLB configuration that uses VRRP. Each pair of ServerIrons provides redundant FWLB, while VRRP on the external pair of ServerIrons provides redundancy for the default gateway address used by the client.

Figure 5.3 Active-Active FWLB with VRRP

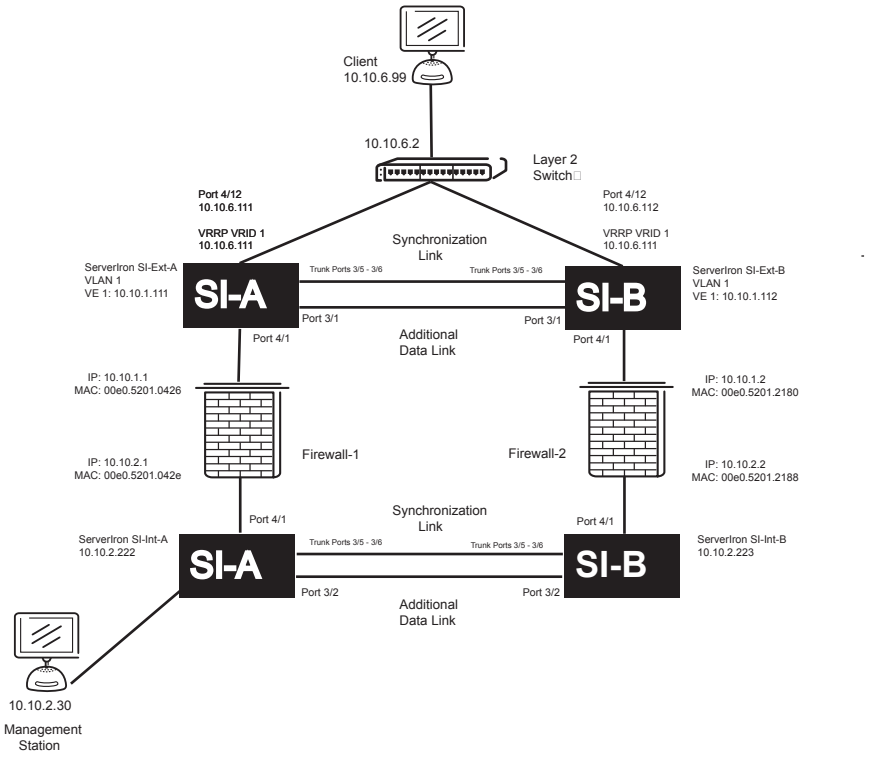

Commands on External ServerIron A (SI-Ext-A)

The following commands change the CLI to the global CONFIG level, then change the hostname to "SI-Ext-A".

```
ServerIron> enable
ServerIron# configure terminal
ServerIron(config)# hostname SI-Ext-A
```

The following commands enable the always-active feature and disable the Spanning Tree Protocol (STP) in VLAN 1, which contains the ports that will carry the FWLB traffic.

```
SI-Ext-A(config)# vlan 1
SI-Ext-A(config-vlan-1)# always-active
SI-Ext-A(config-vlan-1)# no spanning-tree
```

The following commands configure a virtual routing interface on VLAN 1 (the default VLAN), then configure an IP address on the interface. The virtual routing interface is associated with all the ports in the VLAN.

```
SI-Ext-A(config-vlan-1)# router-interface ve 1
SI-Ext-A(config-vlan-1)# exit
SI-Ext-A(config)# interface ve 1
SI-Ext-A(config-ve-1)# ip address 10.10.1.111 255.255.255.0
SI-Ext-A(config-ve-1)# exit
```

The following command configures an IP default route. The next hop for this route is the ServerIron's interface with firewall FW1.

```
SI-Ext-A(config)# ip route 0.0.0.0 0.0.0.0 10.10.1.1
```

The following commands configure port-based VLAN 2, which will contain the port on which VRRP VRID 1 (10.10.6.111) is configured.

```
SI-Ext-A(config)# vlan 2
SI-Ext-A(config-vlan-2)# untag ethernet 4/12
SI-Ext-A(config-vlan-2)# exit
```

The following commands configure the dedicated synchronization link between the ServerIron and its active-active partner. The **trunk** command configures the two ports of the link into a trunk group. The next two commands add the trunk group to a separate port-based VLAN, since the synchronization link must be in its own VLAN. The **server fw-port** command identifies the port number the link is on. If the link is a trunk group, you must specify the MAC address of the group's primary port.

```
SI-Ext-A(config)# trunk switch ethernet 3/5 to 3/6
SI-Ext-A(config)# vlan 10
SI-Ext-A(config-vlan-10)# untagged ethernet 3/5 to 3/6
SI-Ext-A(config-vlan-10)# exit
SI-Ext-A(config)# server fw-port 3/5
```

The following command configures the data link between this ServerIron and its active-active partner. You must use the **server partner-ports** command to specify all the data links with the partner. However, do not use the command for the synchronization link.

NOTE: The **server partner-ports** command is required for all IronClad FWLB configurations in software release 08x.

```
SI-Ext-A(config)# server partner-ports ethernet 3/1
```

The following commands add the firewall definitions. In this example, port HTTP is specified for each firewall. Specifying the application ports on the firewalls is optional. The **port http no-health-check** command under each firewall disables the Layer 4 health check for the HTTP port. When you add an application port to a firewall definition, the ServerIron automatically enables the Layer 4 health check for that port. You must disable the Layer 4 health check if the firewall is unable to act as a proxy for the application and respond to the health check. If the firewall does not respond to the health check, the ServerIron assumes that the port is unavailable and stops sending traffic for the port to the firewall.

The ServerIron will still use a Layer 3 health check (IP ping) to test connectivity to the firewall.

```
SI-Ext-A(config)# server fw-name fw1 10.10.1.1
```



```
SI-Ext-A(config-rs-fw1)# port http
SI-Ext-A(config-rs-fw1)# port http no-health-check
SI-Ext-A(config-rs-fw1)# exit
SI-Ext-A(config)# server fw-name fw2 10.10.1.2
SI-Ext-A(config-rs-fw2)# port http
SI-Ext-A(config-rs-fw2)# port http no-health-check
SI-Ext-A(config-rs-fw2)# exit
```

The following commands add the firewall definitions to the firewall port group (always group 2). The firewall group contains all the ports in VLAN 1 (the default VLAN).

```
SI-Ext-A(config)# server fw-group 2
SI-Ext-A(config-tc-2)# fw-name fw1
SI-Ext-A(config-tc-2)# fw-name fw2
```

The following command enables the active-active mode.

```
SI-Ext-A(config-tc-2)# sym-priority 255
```

NOTE: Do not use the same number on both ServerIrons. For example, use enter **sym-priority 1** on one of the ServerIrons and **sym-priority 255** on the other ServerIron.

The following commands add the paths through the firewalls to the other ServerIron. Each path consists of a path number, a ServerIron port number, the IP address at the other end of the path, and the next-hop IP address. In this example, the topology does not contain routers other than the ServerIrons. If your topology does contain other routers, configure firewall paths for the routers too. For router paths, use the same IP address as the path destination and the next hop.

NOTE: The path IDs must be in contiguous, ascending numerical order, starting with 1. For example, path sequence 1, 2, 3 , 4 is valid. Path sequence 4, 3, 2, 1 or 1, 3, 4, 5 is not valid.

```
SI-Ext-A(config-tc-2)# fwall-info 1 4/1 10.10.2.222 10.10.1.1
SI-Ext-A(config-tc-2)# fwall-info 2 3/1 10.10.2.222 10.10.1.2
SI-Ext-A(config-tc-2)# fwall-info 3 4/1 10.10.2.223 10.10.1.1
SI-Ext-A(config-tc-2)# fwall-info 4 3/1 10.10.2.223 10.10.1.2
```

The following command sets the load balancing method to balance requests based on the firewall that has the least number of connections for the requested service. Since the firewall definitions above specify the HTTP service, the ServerIron will load balance requests based on the firewall that has fewer HTTP session entries in the ServerIron session table.

```
SI-Ext-A(config-tc-2)# fw-predictor per-service-least-conn
```

The following command is part of the always-active feature, which provides the additional data link between the this ServerIron and its partner.

```
SI-Ext-A(config-tc-2)# l2-fwall
SI-Ext-A(config-tc-2)# exit
```

The following commands add static MAC entries for the firewall interfaces with the ServerIron. The static MAC entries are required only if the configuration uses static routes and a single virtual routing interface, as in this example, and if the default gateway for the client or server is the firewall. If the configuration uses a dynamic routing protocol (for example, RIP or OSPF), the static entries are not required. Alternatively, the static entries are not required if you use the ServerIron itself as the default gateway for the client or the server. For example, the static entries are not required if you configure the client to use 10.10.1.111 as its default gateway.

```
SI-Ext-A(config)# vlan 1
SI-Ext-A(config-vlan-1)# static-mac-address 00e0.5201.0426 ethernet 4/1
priority 1 router-type
SI-Ext-A(config-vlan-1)# static-mac-address 00e0.5203.2f80 ethernet 3/1
priority 1 router-type
SI-Ext-A(config-vlan-1)# exit
```

The following commands assign FWLB processing for all forwarding modules to the same WSM CPU. The device uses the same CPU to process all FWLB traffic. You must assign all the traffic to the same WSM CPU. The

commands in this example assign traffic on the forwarding modules in slots 3 and 4 to WSM CPU 1 on the Web Switching Management Module in slot 2.

```
SI-Ext-A(config)# wsm wsm-map slot 3 wsm-slot 2 wsm-cpu 1
SI-Ext-A(config)# wsm wsm-map slot 4 wsm-slot 2 wsm-cpu 1
```

NOTE: For simplicity, the configuration of the other ServerIrons in this example do not include **wsm wsm-map** commands. However, the commands you need to enter depend on the slot locations of the modules in the device and the WSM CPU you want to use.

The following commands enable FWLB.

```
SI-Ext-A(config)# ip l4-policy 1 fw tcp 0 global
SI-Ext-A(config)# ip l4-policy 2 fw udp 0 global
```

The following commands configure the VRRP parameters. The address indicated by the **ip-address** command (10.10.6.111) is the address that will be backed up by VRRP. Since this ServerIron is the owner of the backed up address, the address is configured on the port (this port owns the address) and the address is assigned to the VRID. On external ServerIron B, the VRID will be configured as a backup for 10.10.6.111. The port on which the VRID is configured will have an IP address that is in the same sub-net as the backed up address, but not the same address.

```
ServerIronA(config)# router vrrp
ServerIronA(config)# interface ethernet 4/12
ServerIronA(config-if-4/12)# ip address 10.10.6.111/24
ServerIronA(config-if-4/12)# ip vrrp vrid 1
ServerIronA(config-if-4/12-vrid-1)# owner
ServerIronA(config-if-4/12-vrid-1)# ip-address 10.10.6.111
ServerIronA(config-if-4/12-vrid-1)# activate
ServerIronA(config-if-4/12-vrid-1)# exit
ServerIronA(config-if-4/12)# exit
```

The following command saves the configuration changes to the startup-config file.

```
SI-Ext-A(config)# write memory
```

Commands on External ServerIron B (SI-Ext-B)

Here are the commands for configuring SI-Ext-B. The SLB configuration is identical to the one on SI-Ext-A.

```
ServerIron> enable
ServerIron# configure terminal
ServerIron(config)# hostname SI-Ext-B
SI-Ext-B(config)# vlan 1
SI-Ext-B(config-vlan-1)# always-active
SI-Ext-B(config-vlan-1)# no spanning-tree
SI-Ext-B(config-vlan-1)# router-interface ve 1
SI-Ext-B(config-vlan-1)# exit
SI-Ext-B(config)# interface ve 1
SI-Ext-B(config-ve-1)# ip address 10.10.1.112 255.255.255.0
SI-Ext-B(config-ve-1)# exit
SI-Ext-B(config)# ip route 0.0.0.0 0.0.0.0 10.10.1.1
SI-Ext-B(config)# vlan 2
SI-Ext-B(config-vlan-2)# untag ethernet 4/12
SI-Ext-B(config-vlan-2)# exit
SI-Ext-B(config)# trunk switch ethernet 3/5 to 3/6
SI-Ext-B(config)# vlan 10
SI-Ext-B(config-vlan-10)# untagged ethernet 3/5 to 3/6
SI-Ext-B(config-vlan-10)# exit
SI-Ext-B(config)# server fw-port 3/5
SI-Ext-B(config)# server partner-ports ethernet 3/1
SI-Ext-B(config)# server fw-name fw1 10.10.1.1
SI-Ext-B(config-rs-fw1)# port http
```



```
SI-Ext-B(config-rs-fw1)# port http no-health-check
SI-Ext-B(config-rs-fw1)# exit
SI-Ext-B(config)# server fw-name fw2 10.10.1.2
SI-Ext-B(config-rs-fw2)# port http
SI-Ext-B(config-rs-fw2)# port http no-health-check
SI-Ext-B(config-rs-fw2)# exit
SI-Ext-B(config)# server fw-group 2
SI-Ext-B(config-tc-2)# fw-name fw1
SI-Ext-B(config-tc-2)# fw-name fw2
SI-Ext-B(config-tc-2)# sym-priority 1
SI-Ext-B(config-tc-2)# fwall-info 1 3/1 10.10.2.222 10.10.1.1
SI-Ext-B(config-tc-2)# fwall-info 2 4/1 10.10.2.222 10.10.1.2
SI-Ext-B(config-tc-2)# fwall-info 3 3/1 10.10.2.223 10.10.1.1
SI-Ext-B(config-tc-2)# fwall-info 4 4/1 10.10.2.223 10.10.1.2
SI-Ext-B(config-tc-2)# fw-predictor per-service-least-conn
SI-Ext-B(config-tc-2)# l2-fwall
SI-Ext-B(config-tc-2)# exit
SI-Ext-B(config)# vlan 1
SI-Ext-B(config-vlan-1)# static-mac-address 00e0.5201.0426 ethernet 3/1
priority 1 router-type
SI-Ext-B(config-vlan-1)# static-mac-address 00e0.5203.2f80 ethernet 4/1
priority 1 router-type
SI-Ext-B(config-vlan-1)# exit
SI-Ext-B(config)# ip l4-policy 1 fw tcp 0 global
SI-Ext-B(config)# ip l4-policy 2 fw udp 0 global
ServerIronA(config)# router vrrp
ServerIronA(config)# interface ethernet 4/12
ServerIronA(config-if-4/12)# ip address 10.10.6.112/24
ServerIronA(config-if-4/12)# ip vrrp vrid 1
ServerIronA(config-if-4/12-vrid-1)# backup
ServerIronA(config-if-4/12-vrid-1)# ip-address 10.10.6.111
ServerIronA(config-if-4/12-vrid-1)# activate
ServerIronA(config-if-4/12-vrid-1)# exit
ServerIronA(config-if-4/12)# exit
SI-Ext-B(config)# write memory
```

Commands on Internal ServerIron A (SI-Int-A)

```
ServerIron> enable
ServerIron# configure terminal
ServerIron(config)# hostname SI-Int-A
SI-Int-A(config)# vlan 1
SI-Int-A(config-vlan-1)# always-active
SI-Int-A(config-vlan-1)# no spanning-tree
SI-Int-A(config-vlan-1)# router-interface ve 1
SI-Int-A(config-vlan-1)# exit
SI-Int-A(config)# interface ve 1
SI-Int-A(config-ve-1)# ip address 10.10.2.222 255.255.255.0
SI-Int-A(config-ve-1)# exit
SI-Int-A(config)# ip route 0.0.0.0 0.0.0.0 10.10.2.1
SI-Int-A(config)# trunk switch ethernet 3/5 to 3/6
SI-Int-A(config)# vlan 10
SI-Int-A(config-vlan-10)# untagged ethernet 3/5 to 3/6
SI-Int-A(config-vlan-10)# exit
SI-Int-A(config)# server fw-port 3/5
SI-Int-A(config)# server partner-ports ethernet 3/2
SI-Int-A(config)# server fw-name fw1 10.10.2.1
SI-Int-A(config-rs-fw1)# port http
SI-Int-A(config-rs-fw1)# port http no-health-check
```



```
SI-Int-A(config-rs-fw1)# exit
SI-Int-A(config)# server fw-name fw2 10.10.2.2
SI-Int-A(config-rs-fw2)# port http
SI-Int-A(config-rs-fw2)# port http no-health-check
SI-Int-A(config-rs-fw2)# exit
SI-Int-A(config)# server fw-group 2
SI-Int-A(config-tc-2)# fw-name fw1
SI-Int-A(config-tc-2)# fw-name fw2
SI-Int-A(config-tc-2)# sym-priority 255
SI-Int-A(config-tc-2)# fwall-info 1 4/1 10.10.1.111 10.10.2.1
SI-Int-A(config-tc-2)# fwall-info 2 3/2 10.10.1.111 10.10.2.2
SI-Int-A(config-tc-2)# fwall-info 3 4/1 10.10.1.112 10.10.2.1
SI-Int-A(config-tc-2)# fwall-info 4 3/2 10.10.1.112 10.10.2.2
SI-Int-A(config-tc-2)# fw-predictor per-service-least-conn
SI-Int-A(config-tc-2)# l2-fwall
SI-Int-A(config-tc-2)# exit
SI-Int-A(config)# vlan 1
SI-Int-A(config-vlan-1)# static-mac-address 00e0.5201.042e ethernet 4/1
priority 1 router-type
SI-Int-A(config-vlan-1)# static-mac-address 00e0.5201.2f88 ethernet 3/2
priority 1 router-type
SI-Int-A(config-vlan-1)# exit
SI-Int-A(config)# ip l4-policy 1 fw tcp 0 global
SI-Int-A(config)# ip l4-policy 2 fw udp 0 global
SI-Int-A(config)# write memory
```

Commands on Internal ServerIron B (SI-Int-B)

```
ServerIron> enable
ServerIron# configure terminal
ServerIron(config)# hostname SI-Int-B
SI-Int-B(config)# vlan 1
SI-Int-B(config-vlan-1)# always-active
SI-Int-B(config-vlan-1)# no spanning-tree
SI-Int-B(config-vlan-1)# router-interface ve 1
SI-Int-B(config-vlan-1)# exit
SI-Int-B(config)# interface ve 1
SI-Int-B(config-ve-1)# ip address 10.10.2.223 255.255.255.0
SI-Int-B(config-ve-1)# exit
SI-Int-B(config)# ip route 0.0.0.0 0.0.0.0 10.10.2.2
SI-Int-B(config)# trunk switch ethernet 3/5 to 3/6
SI-Int-B(config)# vlan 10
SI-Int-B(config-vlan-10)# untagged ethernet 3/5 to 3/6
SI-Int-B(config-vlan-10)# exit
SI-Int-B(config)# server fw-port 3/5
SI-Int-B(config)# server partner-ports ethernet 3/2
SI-Int-B(config)# server fw-name fw1 10.10.2.1
SI-Int-B(config-rs-fw1)# port http
SI-Int-B(config-rs-fw1)# port http no-health-check
SI-Int-B(config-rs-fw1)# exit
SI-Int-B(config)# server fw-name fw2 10.10.2.2
SI-Int-B(config-rs-fw2)# port http
SI-Int-B(config-rs-fw2)# port http no-health-check
SI-Int-B(config-rs-fw2)# exit
SI-Int-B(config)# server fw-group 2
SI-Int-B(config-tc-2)# fw-name fw1
SI-Int-B(config-tc-2)# fw-name fw2
SI-Int-B(config-tc-2)# sym-priority 1
SI-Int-B(config-tc-2)# fwall-info 1 3/2 10.10.1.111 10.10.2.1
```



```
SI-Int-B(config-tc-2)# fwall-info 2 4/10 10.10.1.111 10.10.2.2
SI-Int-B(config-tc-2)# fwall-info 3 3/2 10.10.1.112 10.10.2.1
SI-Int-B(config-tc-2)# fwall-info 4 4/10 10.10.1.112 10.10.2.2
SI-Int-B(config-tc-2)# fw-predictor per-service-least-conn
SI-Int-B(config-tc-2)# l2-fwall
SI-Int-B(config-tc-2)# exit
SI-Int-B(config)# vlan 1
SI-Int-B(config-vlan-1)# static-mac-address 00e0.5201.042e ethernet 3/2
priority 1 router-type
SI-Int-B(config-vlan-1)# static-mac-address 00e0.5201.2f88 ethernet 4/1
priority 1 router-type
SI-Int-B(config-vlan-1)# exit
SI-Int-B(config)# ip l4-policy 1 fw tcp 0 global
SI-Int-B(config)# ip l4-policy 2 fw udp 0 global
SI-Int-B(config)# write memory
```


Chapter 6
Configuring Multizone FWLB

Multi-zone FWLB allows you to configure ServerIrons to forward packets based on the destination zone. For example, if your network consists of an Internet side, an internal side, and a Demilitarized Zone (DMZ) in between, you can configure ServerIrons to forward packets through the firewalls to the correct zone.

When you configure multi-zone FWLB, you first identify a zone by configuring standard Access Control Lists (ACLs). An ACL specifies the IP addresses (or address ranges) within the zone. When you configure the firewall group parameters, you add the zones and define them by associating the ACLs with them. Each zone consists of a zone number, an optional name, and a standard ACL that specifies the IP addresses contained in the zone.

You can configure multi-zone FWLB for basic configurations and IronClad (high-availability) configurations. This section provides an example for each type of configuration.

Zone Configuration

When the ServerIron forwards a packet, it selects a path that goes through a firewall to a ServerIron that is in the zone that contains the destination IP address of the packet.

The configuration tasks for multi-zone FWLB are the same as other FWLB implementations, with the exception of configuration for the zones.

When you configure zones:

- Do not define zone 1. When zone 1 is undefined, the zone by default contains all IP addresses that are not explicitly configured as members of other zones (zones 2 – 10). In typical configurations, the ServerIrons in the DMZ and internal network contain zone definitions for each other, while none of the ServerIrons contains a zone definition for zone 1 (thus leaving zone 1 undefined). As a result, traffic that is not destined for an address in the DMZ or the internal network is sent to the Internet.

  You can define zone 1 if you want to, but if you do, this zone contains only the IP address ranges you configure for the zone.
- Do not configure zone information on a ServerIron for the zone the ServerIron is in.
- On the DMZ ServerIron(s), configure zone definitions for the zone(s) in the internal network and other DMZs, if applicable.
- On the internal ServerIron(s), configure zone definitions for the zone(s) in the DMZ(s), and other internal networks, if applicable.

Generally, each ServerIron should contain definitions for two less zones than the total number of zones in the network. The two zones you leave out are zone 1 (which remains undefined) and the zone the ServerIron itself is

in. If you are configuring a ServerIron in zone 1, leave out configuration information for zone 1 and one of the other zones.

Configuring Basic Multi-Zone FWLB

Figure 6.1 shows an example of a basic multi-zone FWLB configuration. In this example, each ServerIron is in a separate zone:

- ServerIron Zone1-SI is in zone 1. By default, zone 1 contains all IP addresses that are not members of other, user-configured zones. You can explicitly configure zone 1 but you do not need to. In the CLI configuration example for this configuration, zone 1 is not configured. ServerIron Zone1-SI contains zone definitions for zone 2 (the DMZ zone) but not for zone 1 or zone 3.
- ServerIron Zone2-SI is in zone 2 (the “DMZ” zone in this example). Zone 2 contains IP addresses in the range 209.157.25.0/24 – 209.157.25.255/24. This ServerIron contains configuration information for zone 3 (the internal network zone) but does not contain definitions for zone 1 (the external network zone) or zone 2 (the DMZ zone itself).
- ServerIron Zone3-SI is in zone 3 (the “internal network” zone in the example). Zone 3 contains IP addresses in the range 209.157.23.0/24 – 209.157.23.255/24. This ServerIron contains configuration information for zone 2 (the DMZ zone) but does not contain definitions for zone 1 (the external network zone) or zone 3 (the internal network zone itself).

When one of the ServerIrons receives traffic whose destination IP address is in another zone, the ServerIron selects a path for the traffic based on the zone the destination IP address is in. For example, if a client on the Internet sends traffic addressed to a server in zone 2, ServerIron Zone1-SI selects a path that sends the traffic through a firewall to ServerIron Zone2-SI, which forwards the traffic to the server. (ServerIron Zone2-SI can be configured to load balance traffic across multiple servers or can simply be used as a Layer 2 switch to forward the traffic to the server.)

When ServerIron Zone2-SI forwards the server’s reply to the client, the ServerIron selects a path to ServerIron Zone1-SI. ServerIron Zone2-SI knows the traffic goes to zone 1 because the destination IP address of the traffic is not in its own sub-net (zone 2) or in zone 3.

Figure 6.1 Basic multi-zone FWLB configuration

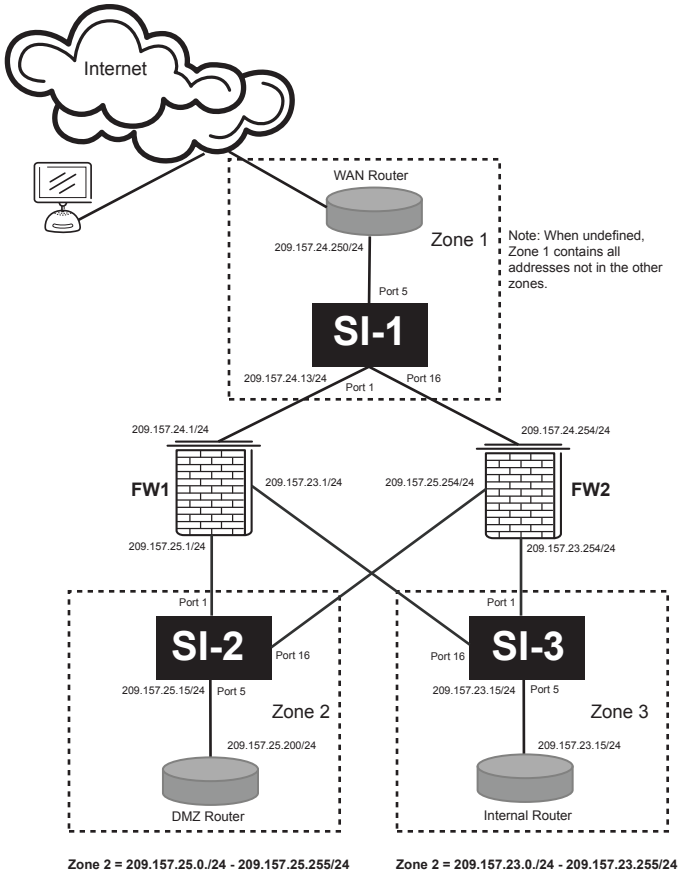

To configure ServerIrons for basic multi-zone FWLB, performs the following tasks:

- **Configure global system parameters.** These parameters include the ServerIron IP address and default gateway. You also need to globally disable the Spanning Tree Protocol (STP). Disabling STP is required for this configuration.
- **Configure global FWLB parameters:**
  - Globally enable FWLB.
  - Identify the port connected to the router.
- **Configure firewall parameters:**
  - Define the firewalls and add them to the firewall group. Each firewall consists of a name and the IP address of its interface with the ServerIron.
- **Configure a standard ACL for each zone the ServerIron is not a member of, except zone 1.** The ACLs identify the IP addresses or address ranges in the other zones. If you leave zone 1 undefined, all IP addresses that are not in this ServerIron's own sub-net and are not members of zones configured on the ServerIron, are assumed to be members of zone 1.

  If the ServerIron is a member of zone 1, configure a standard ACL for all but one of the other zones. In this example, configure an ACL for the DMZ zone (zone 2). The ServerIron will forward traffic that is not addressed to its own sub-net (zone 1) and not addressed to zone 2, to the other zone (zone 3) automatically.
- **Configure firewall group parameters:**
  - Configure the zones. Each zone definition consists of a number, an optional name, and the ACL that specifies the IP addresses in the zone.

  - Configure the paths and add static MAC entries for the firewall interfaces with the ServerIron. Configure a separate path through each firewall to each ServerIron. You also need to configure a path from each ServerIron to the router(s) attached to the ServerIron.

- **Save the configuration to the startup-config file.**

Configuration Example for Basic Multi-Zone FWLB

The following sections show all the ServerIron commands you would enter on each ServerIron to implement the configuration shown in Figure 6.1 on page 6-3.

Most of the configuration tasks for multi-zone FWLB are the same as the tasks for other FWLB configurations. See the other sections in this chapter for procedures.

Commands on ServerIron Zone1-SI

The following commands configure ServerIron "Zone1-SI" in zone 1 in Figure 6.1 on page 6-3.

The first set of commands changes the device name, configures the management IP address, and specifies the default gateway. Notice that the management IP address is in the same sub-net as the firewall interface with the ServerIron. If the ServerIron and the firewall are in different sub-nets, you need to configure source IP addresses and enable source NAT.

In this configuration, the default gateway is the IP address of the one of the firewall interfaces with the ServerIron. In this case, the IP address is the address of firewall FW1's interface with this ServerIron.

```
ServerIron(config)# hostname Zone1-SI
Zone1-SI(config)# ip address 209.157.24.13 255.255.255.0
Zone1-SI(config)# ip default-gateway 209.157.24.1
```

The following command disables the Spanning Tree Protocol (STP). You must disable STP on all the devices in this type of FWLB configuration.

```
Zone1-SI(config)# no span
```

The following commands enable FWLB. Enter the commands exactly as shown for all FWLB configurations. The "0" parameter is required and enables the ServerIron to provide FWLB for all packets of the specified type (TCP or UDP). FWLB is enabled globally. You cannot enable the feature locally, on individual ports.

```
Zone1-SI(config)# ip policy 1 fw tcp 0 global
Zone1-SI(config)# ip policy 2 fw udp 0 global
```

The following command identifies the router port, which is the ServerIron ports connected to a router. In the example in Figure 6.1 on page 6-3, each ServerIron has one router port. If the link is a trunk group, enter the primary port number. In this example, the router port is port 5.

```
Zone1-SI(config)# server router-ports 5
```

The following commands add the firewalls.

```
Zone1-SI(config)# server fw-name FW1 209.157.24.1
Zone1-SI(config-rs-FW1)# exit
Zone1-SI(config)# server fw-name FW2 209.157.24.254
Zone1-SI(config-rs-FW2)# exit
```

The names are specific to the ServerIron and do not need to correspond to any name parameters on the firewalls themselves. The IP addresses are the addresses of the firewall interfaces with the ServerIron.

The following command configures an Access Control List (ACL) for the IP addresses in the DMZ zone (zone 2). The command configures a standard ACL for the addresses in zone 2, which contains addresses in the 209.157.25.x/24 sub-net. The "0.0.0.255" values indicate the significant bits in the IP address you specify. In this case, all bits except the ones in the last node of the address are significant.

In this configuration, only one zone definition is required on each ServerIron, including Zone1-SI. Since the Zone1-SI ServerIron is already in zone 1, the ServerIron will forward packets either to the ServerIron in zone 2 or to the only other ServerIron that is not in zone 2. In this case, the only other ServerIron is the one in zone 3. Thus, if

ServerIron Zone1-SI receives a packet that is not addressed to the sub-net Zone1-SI is in, and is not addressed to a sub-net in zone 2, the ServerIron assumes that the packet is for an address in the other zone, zone 3. The ServerIron forwards the packet to the ServerIron in zone 3.

```
Zone1-SI(config)# access-list 2 permit 209.157.25.0 0.0.0.255
```

Although each zone in this example contains one Class C sub-net, you can configure ACLs for any range of addresses and even for individual host addresses.

NOTE: This example shows a numbered ACL, instead of a named ACL. In the current ServerIron software release, you must use numbered ACLs. The FWLB software does not support zone configuration based on named ACLs.

The following commands configure the firewall group parameters. In this case, the commands configure the firewall zones, add zone 2, and add the firewalls.

```
Zone1-SI(config)# server fw-group 2
Zone1-SI(config-tc-2)# fwall-zone Zone2 2 2
Zone1-SI(config-tc-2)# fw-name FW1
Zone1-SI(config-tc-2)# fw-name FW2
```

The **fwall-zone** command configures a firewall zone. To configure a zone, specify a name for the zone, then a zone number (from 1 – 10), followed by the number of the standard ACL that specifies the IP addresses in the zone. In this example, the ACL number and zone number are the same, but this is not required.

The **fw-name** commands add the firewalls. Specify the names you entered when configuring the firewalls. In this example, the names are “FW1” and “FW2”.

The following commands configure the firewall paths. In the configuration in Figure 6.1 on page 6-3, each ServerIron has five paths:

- A path through FW1 to ServerIron Zone2
- A path through FW2 to ServerIron Zone2
- A path through FW1 to ServerIron Zone3
- A path through FW2 to ServerIron Zone3
- A path to the router

The ServerIron uses the firewall paths to load balance the firewall traffic across the two firewalls. As in other types of FWLB configurations, the paths must be fully meshed among the ServerIrons and firewalls. Thus, the ServerIron has a separate path through each of the firewalls to each of the ServerIrons in the other zones.

The ServerIron also uses the paths for checking the health of the links. The health checking enables the ServerIron to compensate if the link to a firewall becomes unavailable by sending traffic that normally goes through the unavailable firewall through the firewall that is still available.

```
Zone1-SI(config-tc-2)# fwall-info 1 1 209.157.25.15 209.157.24.1
Zone1-SI(config-tc-2)# fwall-info 2 1 209.157.23.11 209.157.24.1
Zone1-SI(config-tc-2)# fwall-info 3 16 209.157.25.15 209.157.24.254
Zone1-SI(config-tc-2)# fwall-info 4 16 209.157.23.11 209.157.24.254
Zone1-SI(config-tc-2)# fwall-info 5 5 209.157.24.250 209.157.24.250
Zone1-SI(config-tc-2)# exit
```

Each **fwall-info** command consists of a path number, a ServerIron port number, the IP address at the other end of the path, and the next-hop IP address. The paths that pass through FW1 use ServerIron port 1, which is connected to FW1. The paths that pass through FW2 use ServerIron port 16.

Notice that the last path, unlike the other paths, has the same IP address for the destination and the next-hop for the path. This path is a router path and ends at the router itself. The other paths are firewall paths and end at the ServerIron at the other end of the firewall.

The following commands add static entries to the ServerIron’s MAC table for the firewall interfaces.

```
Zone1-SI(config)# static-mac-address abcd.5200.348d ethernet 1 high-priority router-
type
```



```
Zone1-SI(config)# static-mac-address abcd.5200.0b50 ethernet 16 high-priority
router-type
```

Each command includes the MAC address of the firewall's interface with the ServerIron and the ServerIron port that is connected to the firewall. The **high-priority** and **router-type** parameters identify the MAC entry type and are required.

NOTE: The syntax for the **static-mac-address** command is slightly different on ServerIron Chassis devices. Instead of a port number, you specify a slot and port number. For the priority, specify **priority 7** instead of **high-priority**.

The following command saves the configuration information to the ServerIron's startup-config file on flash memory. You must save the configuration information before reloading the software or powering down the device. Otherwise, the information is lost.

```
Zone1-SI(config)# write memory
```

Commands on Zone2-SI in Zone 2

The following commands configure ServerIron "Zone2-SI" in zone 2 in Figure 6.1 on page 6-3. The configuration is similar to the one for Zone1-SI, with the following exceptions:

- The management IP address is different.
- The default gateway goes to a different interface on FW1.
- The paths are different due to the ServerIron's placement in the network. (However, like Zone1-SI, ServerIron Zone2-SI has a path through each firewall to the ServerIrons in the other zones, and has a path to its directly attached router.)
- An ACL and zone definition are configured for zone 3. Since this ServerIron is in zone 2, the configuration does not include an ACL and zone definition for zone 2. This ServerIron also does not contain an ACL or zone definition for zone 1. As a result, by default this ServerIron forwards packets that are not addressed to the ServerIron's own sub-net or to a sub-net in zone 3, to zone 1.

```
ServerIron(config)# hostname Zone2-SI
Zone2-SI(config)# ip address 209.157.24.15 255.255.255.0
Zone2-SI(config)# ip default-gateway 209.157.25.1
Zone2-SI(config)# ip policy 1 fw tcp 0 global
Zone2-SI(config)# ip policy 2 fw udp 0 global
Zone2-SI(config)# no span
Zone2-SI(config)# server router-ports 5
Zone2-SI(config)# server fw-name FW1 209.157.25.1
Zone2-SI(config-rs-FW1)# exit
Zone2-SI(config)# server fw-name FW2 209.157.25.254
Zone2-SI(config-rs-FW2)# exit
Zone2-SI(config)# access-list 3 permit 209.157.23.0 0.0.0.255
Zone2-SI(config)# server fw-group 2
Zone2-SI(config-tc-2)# fwall-zone Zone3 3 3
Zone2-SI(config-tc-2)# fw-name FW1
Zone2-SI(config-tc-2)# fw-name FW2
Zone2-SI(config-tc-2)# fwall-info 1 1 209.157.25.15 209.157.24.1
Zone2-SI(config-tc-2)# fwall-info 2 16 209.157.23.11 209.157.24.1
Zone2-SI(config-tc-2)# fwall-info 3 16 209.157.25.15 209.157.24.254
Zone2-SI(config-tc-2)# fwall-info 4 1 209.157.23.11 209.157.24.254
Zone2-SI(config-tc-2)# fwall-info 5 5 209.157.25.200 209.157.25.200
Zone2-SI(config-tc-2)# exit
Zone2-SI(config)# static-mac-address abcd.5200.348b ethernet 1 high-priority router-
type
Zone2-SI(config)# static-mac-address abcd.5200.0b4e ethernet 16 high-priority
router-type
Zone2-SI(config)# write memory
```



```
Zone2-SI(config)# exit
Zone2-SI# reload
```

Commands on Zone3-SI in Zone 3

The following commands configure ServerIron “Zone3-SI” in zone 3 in Figure 6.2 on page 6-8. The configuration is similar to the ones for the other ServerIrons, with the following exceptions:

- The management IP address is different.
- The default gateway goes to an interface on FW2.
- The paths are different due to the ServerIron’s placement in the network.
- An ACL and zone definition are configured for zone 2. Since this ServerIron is in zone 3, the configuration does not include an ACL and zone definition for the zone. This ServerIron also does not contain an ACL or zone definition for zone 1. As a result, by default this ServerIron forwards packets that are not addressed to the ServerIron’s own sub-net or to a sub-net in zone 2, to zone 1.

```
ServerIron(config)# hostname Zone3-SI
Zone3-SI(config)# ip address 209.157.23.11 255.255.255.0
Zone3-SI(config)# ip default-gateway 209.157.23.1
Zone3-SI(config)# no span
Zone3-SI(config)# ip policy 1 fw tcp 0 global
Zone3-SI(config)# ip policy 2 fw udp 0 global
Zone3-SI(config)# server router-ports 5
Zone3-SI(config)# server fw-name FW1 209.157.23.1
Zone3-SI(config-rs-FW1)# exit
Zone3-SI(config)# server fw-name FW2 209.157.23.254
Zone3-SI(config-rs-FW2)# exit
Zone3-SI(config)# access-list 2 permit 209.157.25.0 0.0.0.255
Zone3-SI(config)# server fw-group 2
Zone3-SI(config-tc-2)# fwall-zone Zone2 2 2
Zone3-SI(config-tc-2)# fw-name FW1
Zone3-SI(config-tc-2)# fw-name FW2
Zone3-SI(config-tc-2)# fwall-info 1 16 209.157.24.13 209.157.23.1
Zone3-SI(config-tc-2)# fwall-info 2 1 209.157.24.13 209.157.23.254
Zone3-SI(config-tc-2)# fwall-info 3 16 209.157.25.15 209.157.23.1
Zone3-SI(config-tc-2)# fwall-info 4 1 209.157.25.15 209.157.23.254
Zone3-SI(config-tc-2)# fwall-info 5 5 209.157.23.15 209.157.23.15
Zone3-SI(config-tc-2)# exit
Zone3-SI(config)# static-mac-address abcd.5200.3489 ethernet 16 high-priority
router-type
Zone3-SI(config)# static-mac-address abcd.5200.0b4c ethernet 1 high-priority router-
type
Zone3-SI(config)# write memory
Zone3-SI(config)# exit
Zone3-SI# reload
```

Configuring IronClad Multi-Zone FWLB

Figure 6.2 on page 6-8 shows an example of an IronClad (high-availability) multi-zone FWLB configuration. This example has the same zones as the basic example in Figure 6.1 on page 6-3, but in the IronClad configuration each zone contains a pair of active-standby ServerIrons instead of a single ServerIron.

In this configuration, the ServerIrons on the left side of Figure 6.1 are the active ServerIrons. The ServerIrons on the right are the standby ServerIrons. Each active-standby pair is connected by a private link, which the ServerIrons use to exchange failover information. The ports used by the private links are in their own port-based VLAN, separate from the other ServerIron ports. Add the ports as untagged ports. For added redundancy, the private links also are configured as two-port trunk groups.

This example also uses a simplified topology. Instead of using Layer 2 switches and redundant links to provide failover data paths from the devices on the left side to the devices on the right side, this configuration uses additional links between the ServerIrons. The L2-fwall and always-active options enable you to use this type of simplified topology. The ***L2-fwall*** option prevents data loops by blocking traffic on the standby ServerIron, while the ***always-active*** option allows the standby ServerIrons to pass traffic to their active partners for forwarding.

Figure 6.2 High-availability configuration with separate firewall zones

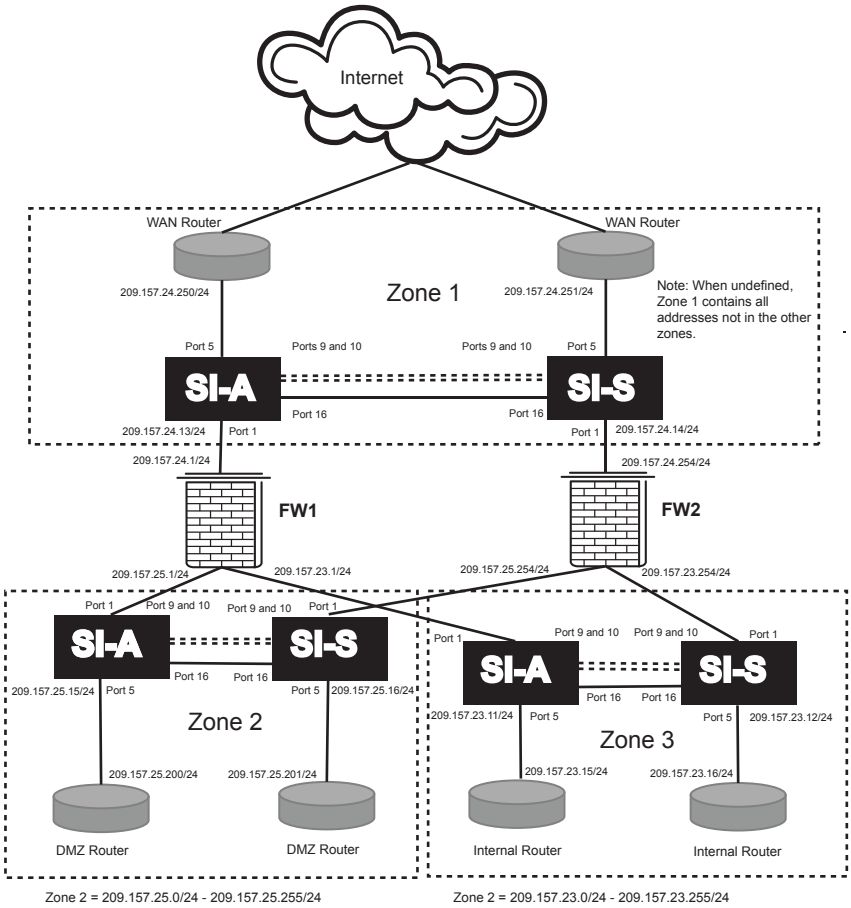

To configure ServerIrons for IronClad multi-zone FWLB, performs the following tasks:

- **Configure global system parameters.** These parameters include the ServerIron IP address and default gateway. You also need to globally disable the Spanning Tree Protocol (STP). Disabling STP is required for this configuration.
- **Configure global FWLB parameters:**
  - Globally enable FWLB.
  - Identify the synchronization port, which is the port connected to this ServerIron's high-availability partner and place the port in a separate Layer port-based VLAN, as an untagged port. (This task applies only to high-availability configurations.)
  - Identify the port connected to the router.
  - Enable the always-active feature for the VLAN that contains all the ports except the synchronization link.

- **Configure a standard ACL for each zone the ServerIron is not a member of, except zone 1.** The ACLs identify the IP addresses or address ranges in the other zones. If you leave zone 1 undefined, all IP addresses that are not in this ServerIron's own sub-net and are not members of zones configured on the ServerIron, are assumed to be members of zone 1.

  If the ServerIron is a member of zone 1, configure a standard ACL for all but one of the other zones. In this example, configure an ACL for the DMZ zone (zone 3). The ServerIron will forward traffic that is not addressed to its own sub-net and not addressed to zone 2, to the other zone (zone 3) automatically.

- **Configure firewall parameters:**
  - Define the firewalls and add them to the firewall group. Each firewall consists of a name and the IP address of its interface with the ServerIron.
- **Configure firewall group parameters:**
  - Configure the zones. Each zone definition consists of a number, an optional name, and the ACL that specifies the IP addresses in the zone.
  - Configure the paths and add static MAC entries for the firewall interfaces with the ServerIron. Configure a separate path through each firewall to each ServerIron. You also need to configure a path from each ServerIron to the router(s) attached to the ServerIron.
  - Specify the ServerIron priority. The ServerIron with the higher priority value is the ServerIron in the active-standby pair that is active by default.
- **Save the configuration to the startup-config file.**
- **Reload the software.** This step is required to place the trunk groups into effect.

Failover Algorithm

ServerIrons in high-availability multi-zone FWLB configurations use the following criteria for failover:

- Connection to zones – If one ServerIron in an active-standby ServerIron has connectivity to more zones than the other ServerIron, the ServerIron with connectivity to more zones is the active ServerIron.
- Total number of good paths – If each ServerIron has connectivity to an equal number of zones, the ServerIron with more good paths, within the configured tolerance, is the active ServerIron. The paths include firewall paths and router paths. By default, the ServerIrons can tolerate up to half of the firewall paths and half the router paths being down before failover based on good paths occurs. You can change the path tolerance.
- Priority – If all the above metrics are equal on each ServerIron, the ServerIron with the higher priority is the active ServerIron.

Configuration Example for IronClad Multi-Zone FWLB

The following sections show all the ServerIron commands you would enter on each ServerIron to implement the configuration shown in Figure 6.2 on page 6-8.

Most of the configuration tasks for multi-zone FWLB are the same as the tasks for other FWLB configurations. See the other sections in this chapter for procedures.

Commands on Zone1-SI-A Zone 1

The following commands configure ServerIron "Zone1-SI-A", on the left side of the zone 1 in Figure 6.2 on page 6-8.

The following commands change the device name, configure the management IP address, and specify the default gateway. Notice that the management IP address is in the same sub-net as the firewall interface with the ServerIron. If the ServerIron and the firewall are in different sub-nets, you need to configure source IP addresses and enable source NAT.

In this configuration, the default gateway for each ServerIron is the IP address of the firewall interface with that ServerIron. In this case, the IP address is the address of firewall FW1's interface with this ServerIron.

```
ServerIron(config)# hostname Zone1-SI-A
Zone1-SI-A(config)# ip address 209.157.24.13 255.255.255.0
Zone1-SI-A(config)# ip default-gateway 209.157.24.1
```

The following command disables the Spanning Tree Protocol (STP). You must disable STP on all the devices in this type of FWLB configuration.

```
Zone1-SI-A(config)# no span
```

The following commands enable FWLB. Enter the commands exactly as shown for all FWLB configurations. The "0" parameter is required and enables the ServerIron to provide FWLB for all packets of the specified type (TCP or UDP). FWLB is enabled globally. You cannot enable the feature locally, on individual ports.

```
Zone1-SI-A(config)# ip policy 1 fw tcp 0 global
Zone1-SI-A(config)# ip policy 2 fw udp 0 global
```

The following command identifies the router port, which is the ServerIron port connected to a router. In the example in Figure 6.2 on page 6-8, each ServerIron has one router port.

```
Zone1-SI-A(config)# server router-ports 5
```

The following commands identify the port for the link to the other ServerIron. If the link is a trunk group, enter the primary port number. In this example, the link is a trunk group made of ports 9 and 10, but you only need to specify port 9, the trunk group's primary port.

The commands also create a trunk group for the ports that connect this ServerIron to its high-availability partner, then create a separate port-based VLAN containing the ports in the trunk group. Always configure the private link between the active and standby ServerIron in a separate port-based VLAN. Add the ports as untagged ports.

Using a trunk group for the link between the active and standby ServerIrons is not required, but using a trunk group adds an additional level of redundancy for enhanced availability. If one of the ports in a trunk group goes down, the link remains intact as long as the other port remains up. Make sure you configure a server trunk group, not a switch trunk group. The default trunk group type is switch, so you must specify the **server** option. Trunk groups require a software reload to take effect, so after you complete the ServerIron configuration and the save the configuration to flash memory, you need to reload the software.

Notice that the **server fw-port** command (which identifies the port connected to the other ServerIron) refers to only one port, even though the link is actually a multiple-port trunk group. This port number is the primary port of the trunk group. If you use a trunk group for the private link between the active and standby ServerIrons, refer to the group by its primary port, in this case port 9.

```
Zone1-SI-A(config)# server fw-port 9
Zone1-SI-A(config)# trunk server ethernet 9 to 10
Zone1-SI-A(config)# vlan 10 by port
Zone1-SI-A(config-vlan-10)# untagged 9 to 10
Zone1-SI-A(config-vlan-10)# exit
```

The following commands enable the always-active option on the default VLAN.

The default VLAN contains all the ports you have not placed in other port-based VLANs. In this configuration, the default VLAN contains all ports except ports 9 and 10, which are used for the private link between the active and standby ServerIrons.

The always-active option enables the standby ServerIron to forward traffic by sending it through the active ServerIron. This option is useful in configurations where you need to enable the L2-fwall option (to prevent Layer 2 loops through the standby ServerIron), but you also need to allow traffic to pass through the standby ServerIron because that ServerIron is the only path for some traffic.

Without the always-active option, the standby ServerIron blocks all traffic. As a result, if the router connected to the standby ServerIron forwards client traffic addressed to a server in the DMZ, the traffic is blocked by the standby ServerIron. However, when the always-active option is enabled, the standby ServerIron forwards traffic to its active partner ServerIron, which then forwards the traffic to its destination.

In some configurations, you do not need the L2-fwall option or the always-active option. However, configurations that do not use these options compensate with redundant links and sometimes extra Layer 2 switches. For

example, if each ServerIron in Figure 6.2 on page 6-8 had links to both routers in its zone and also to both firewalls, and if Layer 2 switches were added to the configuration to allow STP to prevent Layer 2 loops, then it is possible that neither the l2-fwall nor the always-active option would be required.

In the configuration in Figure 6.2 on page 6-8, each router and firewall is connected to only one of the two ServerIrons in an active-standby pair. Neither the routers nor the firewalls have direct links (or links through Layer 2 switches) to both the active and standby ServerIrons in their zones.

Using the L2-fwall and always-active options allows you to simplify the network topology while still obtaining the benefits of the IronClad (high-availability) configuration. Use the following commands to enable the always-active option in the default VLAN (VLAN 1). You enable the L2-fwall option when you configure firewall group parameters (see below).

```
Zone1-SI-A(config)# vlan 1
Zone1-SI-A(config-vlan-1)# always-active
Zone1-SI-A(config-vlan-1)# exit
```

The following commands add the firewalls.

```
Zone1-SI-A(config)# server fw-name FW1 209.157.24.1
Zone1-SI-A(config-rs-FW1)# exit
Zone1-SI-A(config)# server fw-name FW2 209.157.24.254
Zone1-SI-A(config-rs-FW2)# exit
```

The names are specific to the ServerIron and do not need to correspond to any name parameters on the firewalls themselves. The IP addresses are the addresses of the firewall interfaces with the ServerIron.

The following command configures an Access Control List (ACL) for the IP addresses in one of the zones this ServerIron is not in. In this configuration, only one zone definition is required on each ServerIron, including Zone1-SI-A and Zone1-SI-S. Since the active Zone 1 ServerIron is already in zone 1, the ServerIron will forward packets either to the active ServerIron in zone 2 or to the only other active ServerIron that is not in zone 2. In this case, that other active ServerIron is in zone 3. Thus, if ServerIron Zone1-SI-A receives a packet that is not addressed to the sub-net Zone1-SI-A is in, and is not addressed to a sub-net in zone 2, the ServerIron assumes that the packet is for an address in the other zone, zone 3. The ServerIron forwards the packet to the ServerIron in zone 3.

The command configures an ACL for the addresses in zone 2, which contains addresses in the 209.157.25.x/24 sub-net. The “0.0.0.255” values indicate the significant bits in the IP address you specify. In this case, all bits except the ones in the last node of the address are significant.

```
Zone1-SI-A(config)# access-list 2 permit 209.157.25.0 0.0.0.255
```

Although each zone in this example contains one Class C sub-net, you can configure ACLs for any range of addresses and even for individual host addresses.

NOTE: This example shows a numbered ACL, instead of a named ACL. In the current ServerIron software release, you must use numbered ACLs. The FWLB software does not support zone configuration based on named ACLs.

The following commands configure the firewall group parameters. In this case, the commands configure the firewall zones, add the firewalls, enable the L2-fwall option, and set the active-standby priority.

```
Zone1-SI-A(config)# server fw-group 2
Zone1-SI-A(config-tc-2)# fwall-zone Zone2 2 2
Zone1-SI-A(config-tc-2)# fw-name FW1
Zone1-SI-A(config-tc-2)# fw-name FW2
Zone1-SI-A(config-tc-2)# l2-fwall
Zone1-SI-A(config-tc-2)# sym-priority 255
```

The **fwall-zone** command configures a firewall zone. To configure a zone, specify a name for the zone, then a zone number (from 1 – 10), followed by the number of the standard ACL that specifies the IP addresses in the zone. In this example, the ACL numbers and zone numbers are the same, but this is not required.

The **fw-name** commands add the firewalls. Specify the names you entered when configuring the firewalls. In this example, the names are “FW1” and “FW2”.

The **l2-fwall** command enables the L2-fwall option. This option blocks the Layer 2 traffic on the standby ServerIrons. If you do not enable this mode, Layer 2 traffic can pass through the ServerIrons, causing loops. Layer 3 traffic is automatically blocked on the standby ServerIrons, so you do not need to explicitly block the traffic. The always-active option (enabled in the default VLAN in commands described earlier) allows the standby ServerIron to still forward traffic by sending the traffic to the active ServerIron over the private link between the ServerIrons.

The **sym-priority** command specifies the priority of this ServerIron with respect to the other ServerIron for the firewalls in the firewall group. The priority can be from 0 – 255. The ServerIron with the higher priority is the default active ServerIron for the firewalls within the group.

NOTE: If you specify 0, the CLI removes the priority. When you save the configuration to the startup-config file, the **sym-priority** command is removed. Use this method to remove the priority. You cannot remove the priority using the **no sym-priority** command.

The following commands configure the firewall paths. In the configuration in Figure 6.2 on page 6-8, each ServerIron has nine paths:

- A path through FW1 to ServerIron Zone3-SI-A, the active ServerIron in zone 3.
- A path through FW2 to ServerIron Zone3-SI-A. (This path passes through the standby ServerIron, then through FW2.)
- A path through FW1 to ServerIron Zone3-SI-S, the standby ServerIron in zone 3.
- A path through FW2 to ServerIron Zone3-SI-S. (This path passes through the standby ServerIron.)
- A path through FW1 to ServerIron Zone2-SI-A.
- A path through FW2 to ServerIron Zone2-SI-A.
- A path through FW1 to ServerIron Zone2-SI-S.
- A path through FW2 to ServerIron Zone2-SI-S.
- A path to the router.

The ServerIron uses the firewall paths to load balance the firewall traffic across the two firewalls. As in other types of FWLB configurations, the paths must be fully meshed among the ServerIrons and firewalls. Thus, the ServerIron has a separate path through each of the firewalls to each of the ServerIrons in the other zones.

The ServerIron also uses the paths for checking the health of the links. The health checking enables the ServerIron to compensate if the link to a firewall becomes unavailable by sending traffic that normally goes through the unavailable firewall through the firewall that is still available. The results of the path health checks also play a role in the failover mechanism. The ServerIron determines how many zones it can access and how many firewall and router paths are good based on health checks of the paths. If a path fails a health check, this can result in a failover to the other ServerIron. (See “Failover Algorithm” on page 6-9.)

```
Zone1-SI-A(config-tc-2)# fwall-info 1 1 209.157.23.11 209.157.24.1
Zone1-SI-A(config-tc-2)# fwall-info 2 1 209.157.23.12 209.157.24.1
Zone1-SI-A(config-tc-2)# fwall-info 3 16 209.157.23.11 209.157.24.254
Zone1-SI-A(config-tc-2)# fwall-info 4 16 209.157.23.12 209.157.24.254
Zone1-SI-A(config-tc-2)# fwall-info 5 1 209.157.25.15 209.157.24.1
Zone1-SI-A(config-tc-2)# fwall-info 6 1 209.157.25.16 209.157.24.1
Zone1-SI-A(config-tc-2)# fwall-info 7 16 209.157.25.15 209.157.24.254
Zone1-SI-A(config-tc-2)# fwall-info 8 16 209.157.25.16 209.157.24.254
Zone1-SI-A(config-tc-2)# fwall-info 9 5 209.157.24.250 209.157.24.250
Zone1-SI-A(config-tc-2)# exit
```

Each **fwall-info** command consists of a path number, a ServerIron port number, the IP address at the other end of the path, and the next-hop IP address. The paths that pass through FW1 use ServerIron port 1, which is connected to FW1. The paths that pass through FW2 (by way of the standby ServerIron, Zone1-SI-S) use ServerIron port 16, which is connected to Zone1-SI-S. Note that the connection on port 16 is different from the private link between the two ServerIrons on ports 9 and 10. The connection on port 16 is in the same VLAN as the links to the routers and firewalls (the default VLAN, VLAN 1). The private link on ports 9 and 10 is in a separate

port-based VLAN and is not used in any of the paths. The private link on ports 9 and 10 in VLAN 2 is used only to exchange failover information. All traffic between zones uses the links in the default VLAN.

Notice that the last path, unlike the other paths, has the same IP address for the destination and the next-hop for the path. This path is a router path and ends at the router itself. The other paths are firewall paths and end at the ServerIron at the other end of the firewall.

The following commands add static entries to the ServerIron's MAC table for the firewall interfaces.

```
Zone1-SI-A(config)# vlan 1
Zone1-SI-A(config-vlan-1)# static-mac-address abcd.5200.348d ethernet 1 high-
priority router-type
Zone1-SI-A(config-vlan-1)# static-mac-address abcd.5200.0b50 ethernet 16 high-
priority router-type
Zone1-SI-A(config-vlan-1)# exit
```

Each command includes the MAC address of the firewall's interface with the ServerIron and the ServerIron port that is connected to the firewall. The **high-priority** and **router-type** parameters identify the MAC entry type and are required.

NOTE: If you enter the command at the global CONFIG level, the static MAC entry applies to the default port-based VLAN (VLAN 1). If you enter the command at the configuration level for a specific port-based VLAN, the entry applies to that VLAN and not to the default VLAN.

NOTE: The syntax for the **static-mac-address** command is slightly different on ServerIron Chassis devices. Instead of a port number, you specify a slot and port number. For the priority, specify **priority 7** instead of **high-priority**.

The following command saves the configuration information to the ServerIron's startup-config file on flash memory. You must save the configuration information before reloading the software or powering down the device. Otherwise, the information is lost.

```
Zone1-SI-A(config)# write memory
```

The following commands change the CLI to the Privileged EXEC level, and reload the software. Since this configuration includes a trunk group, you must reload the software to place the trunk group into effect.

```
Zone1-SI-A(config)# exit
Zone1-SI-A# reload
```

Commands on Zone1-SI-S in Zone 1

The following commands configure ServerIron "Zone1-SI-S", on the right side of zone 1 in Figure 6.2 on page 6-8. The configuration is similar to the one for Zone1-SI-A, with the following exceptions:

- The management IP address is different.
- The default gateway goes to firewall FW2's interface with the ServerIron. (The default gateway for Zone1-SI-A goes to FW1's interface with that ServerIron.)
- The priority is set to 1 instead of 255. The lower priority makes this ServerIron the standby ServerIron by default.
- The paths are different due to the ServerIron's placement in the network. (However, like Zone1-SI-A, ServerIron Zone1-SI-S has a path through each firewall to each of the ServerIrons in the other zones, and has a path to its directly attached router.)

```
ServerIron(config)# hostname Zone1-SI-S
Zone1-SI-S(config)# ip address 209.157.24.14 255.255.255.0
Zone1-SI-S(config)# ip default-gateway 209.157.24.254
Zone1-SI-S(config)# no span
Zone1-SI-S(config)# ip policy 1 fw tcp 0 global
Zone1-SI-S(config)# ip policy 2 fw udp 0 global
Zone1-SI-S(config)# server router-ports 5
```



```
Zone1-SI-S(config)# server fw-port 9
Zone1-SI-S(config)# trunk switch ethernet 9 to 10
Zone1-SI-S(config)# vlan 10 by port
Zone1-SI-S(config-vlan-10)# untagged 9 to 10
Zone1-SI-S(config-vlan-10)# exit
Zone1-SI-S(config)# vlan 1
Zone1-SI-S(config-vlan-1)# always-active
Zone1-SI-S(config-vlan-1)# exit
Zone1-SI-S(config)# server fw-name FW1 209.157.24.1
Zone1-SI-S(config-rs-FW1)# exit
Zone1-SI-S(config)# server fw-name FW2 209.157.24.254
Zone1-SI-S(config-rs-FW2)# exit
Zone1-SI-S(config)# access-list 2 permit 209.157.25.0 0.0.0.255
Zone1-SI-S(config)# server fw-group 2
Zone1-SI-S(config-tc-2)# fwall-zone Zone2 2 2
Zone1-SI-S(config-tc-2)# fw-name FW1
Zone1-SI-S(config-tc-2)# fw-name FW2
Zone1-SI-S(config-tc-2)# l2-fwall
Zone1-SI-S(config-tc-2)# sym-priority 1
Zone1-SI-S(config-tc-2)# fwall-info 1 16 209.157.23.11 209.157.24.1
Zone1-SI-S(config-tc-2)# fwall-info 2 16 209.157.23.12 209.157.24.1
Zone1-SI-S(config-tc-2)# fwall-info 3 1 209.157.23.11 209.157.24.254
Zone1-SI-S(config-tc-2)# fwall-info 4 1 209.157.23.12 209.157.24.254
Zone1-SI-S(config-tc-2)# fwall-info 5 16 209.157.25.15 209.157.24.1
Zone1-SI-S(config-tc-2)# fwall-info 6 16 209.157.25.16 209.157.24.1
Zone1-SI-S(config-tc-2)# fwall-info 7 1 209.157.25.15 209.157.24.254
Zone1-SI-S(config-tc-2)# fwall-info 8 1 209.157.25.16 209.157.24.254
Zone1-SI-S(config-tc-2)# fwall-info 9 5 209.157.24.251 209.157.24.251
Zone1-SI-S(config-tc-2)# exit
Zone1-SI-S(config)# vlan 1
Zone1-SI-S(config-vlan-1)# static-mac-address abcd.5200.348d ethernet 1 high-
priority router-type
Zone1-SI-S(config-vlan-1)# static-mac-address abcd.5200.0b50 ethernet 16 high-
priority router-type
Zone1-SI-S(config-vlan-1)# exit
Zone1-SI-S(config)# write memory
Zone1-SI-S(config)# exit
Zone1-SI-S# reload
```

Commands on Zone2-SI-A in Zone 2

The following commands configure ServerIron “Zone2-SI-A”, on the left side of zone 2 in Figure 6.2 on page 6-8. The configuration is similar to the one for the active ServerIron in zone 1, with the following exceptions:

- The management IP address is different.
- The default gateway goes to a different interface on FW1.
- The paths are different due to the ServerIron’s placement in the network. (However, like Zone1-SI-A and Zone1-SI-S, ServerIron Zone1-SI-S has a path through each firewall to each of the ServerIrons in the other zones, and has a path to its directly attached router.)
- Only one ACL and zone definition are configured, for zone 3. Since this ServerIron is in zone 2, the configuration does not include an ACL and zone definition for the zone. This ServerIron also does not contain an ACL or zone definition for zone 1. As a result, by default this ServerIron forwards packets that are not addressed to the ServerIron’s own sub-net or to a sub-net in zone 3, to zone 1.

```
ServerIron(config)# hostname Zone2-SI-A
Zone2-SI-A(config)# ip address 209.157.24.15 255.255.255.0
Zone2-SI-A(config)# ip default-gateway 209.157.25.1
Zone2-SI-A(config)# no span
```



```
Zone2-SI-A(config)# ip policy 1 fw tcp 0 global
Zone2-SI-A(config)# ip policy 2 fw udp 0 global
Zone2-SI-A(config)# server router-ports 5
Zone2-SI-A(config)# server fw-port 9
Zone2-SI-A(config)# trunk switch ethernet 9 to 10
Zone2-SI-A(config)# vlan 10 by port
Zone2-SI-A(config-vlan-10)# untagged 9 to 10
Zone2-SI-A(config-vlan-10)# exit
Zone2-SI-A(config)# vlan 1
Zone2-SI-A(config-vlan-1)# always-active
Zone2-SI-A(config-vlan-1)# exit
Zone2-SI-A(config)# server fw-name FW1 209.157.25.1
Zone2-SI-A(config-rs-FW1)# exit
Zone2-SI-A(config)# server fw-name FW2 209.157.25.254
Zone2-SI-A(config-rs-FW2)# exit
Zone2-SI-A(config)# access-list 3 permit 209.157.23.0 0.0.0.255
Zone2-SI-A(config)# server fw-group 2
Zone2-SI-A(config-tc-2)# fwall-zone Zone3 3 3
Zone2-SI-A(config-tc-2)# fw-name FW1
Zone2-SI-A(config-tc-2)# fw-name FW2
Zone2-SI-A(config-tc-2)# l2-fwall
Zone2-SI-A(config-tc-2)# sym-priority 1
Zone2-SI-A(config-tc-2)# fwall-info 1 1 209.157.23.11 209.157.25.1
Zone2-SI-A(config-tc-2)# fwall-info 2 1 209.157.23.12 209.157.25.1
Zone2-SI-A(config-tc-2)# fwall-info 3 1 209.157.24.13 209.157.25.1
Zone2-SI-A(config-tc-2)# fwall-info 4 1 209.157.24.14 209.157.25.1
Zone2-SI-A(config-tc-2)# fwall-info 5 16 209.157.23.11 209.157.25.254
Zone2-SI-A(config-tc-2)# fwall-info 6 16 209.157.23.12 209.157.25.254
Zone2-SI-A(config-tc-2)# fwall-info 7 16 209.157.24.13 209.157.25.254
Zone2-SI-A(config-tc-2)# fwall-info 8 16 209.157.24.14 209.157.25.254
Zone2-SI-A(config-tc-2)# fwall-info 9 5 209.157.25.200 209.157.25.200
Zone2-SI-A(config-tc-2)# exit
Zone2-SI-A(config)# vlan 1
Zone2-SI-A(config-vlan-1)# static-mac-address abcd.5200.348b ethernet 1 high-
priority router-type
Zone2-SI-A(config-vlan-1)# static-mac-address abcd.5200.0b4e ethernet 16 high-
priority router-type
Zone2-SI-A(config-vlan-1)# exit
Zone2-SI-A(config)# write memory
Zone2-SI-A(config)# exit
Zone2-SI-A# reload
```

Commands on Zone2-SI-S in Zone 2

The following commands configure ServerIron “Zone2-SI-S”, on the right side of zone 2 in Figure 6.2 on page 6-8.

```
ServerIron(config)# hostname Zone2-SI-S
Zone2-SI-S(config)# ip address 209.157.25.16 255.255.255.0
Zone2-SI-S(config)# ip default-gateway 209.157.25.254
Zone2-SI-S(config)# no span
Zone2-SI-S(config)# ip policy 1 fw tcp 0 global
Zone2-SI-S(config)# ip policy 2 fw udp 0 global
Zone2-SI-S(config)# server router-ports 5
Zone2-SI-S(config)# server fw-port 9
Zone2-SI-S(config)# trunk switch ethernet 9 to 10
Zone2-SI-S(config)# vlan 10 by port
Zone2-SI-S(config-vlan-10)# untagged 9 to 10
Zone2-SI-S(config-vlan-10)# exit
Zone2-SI-S(config)# vlan 1
```



```
Zone2-SI-S(config-vlan-1)# always-active
Zone2-SI-S(config-vlan-1)# exit
Zone2-SI-S(config)# server fw-name FW1 209.157.25.1
Zone2-SI-S(config-rs-FW1)# exit
Zone2-SI-S(config)# server fw-name FW2 209.157.25.254
Zone2-SI-S(config-rs-FW2)# exit
Zone2-SI-S(config)# access-list 3 permit 209.157.23.0 0.0.0.255
Zone2-SI-S(config)# server fw-group 2
Zone2-SI-S(config-tc-2)# fwall-zone Zone3 3 3
Zone2-SI-S(config-tc-2)# fw-name FW1
Zone2-SI-S(config-tc-2)# fw-name FW2
Zone2-SI-S(config-tc-2)# l2-fwall
Zone2-SI-S(config-tc-2)# sym-priority 1
Zone2-SI-S(config-tc-2)# fwall-info 1 16 209.157.23.11 209.157.25.1
Zone2-SI-S(config-tc-2)# fwall-info 2 16 209.157.23.12 209.157.25.1
Zone2-SI-S(config-tc-2)# fwall-info 3 16 209.157.24.13 209.157.25.1
Zone2-SI-S(config-tc-2)# fwall-info 4 16 209.157.24.14 209.157.25.1
Zone2-SI-S(config-tc-2)# fwall-info 5 1 209.157.23.11 209.157.25.254
Zone2-SI-S(config-tc-2)# fwall-info 6 1 209.157.23.12 209.157.25.254
Zone2-SI-S(config-tc-2)# fwall-info 7 1 209.157.24.13 209.157.25.254
Zone2-SI-S(config-tc-2)# fwall-info 8 1 209.157.24.14 209.157.25.254
Zone2-SI-S(config-tc-2)# fwall-info 9 5 209.157.25.200 209.157.25.201
Zone2-SI-S(config-tc-2)# exit
Zone2-SI-S(config)# vlan 1
Zone2-SI-S(config-vlan-1)# static-mac-address abcd.5200.348b ethernet 1 high-
priority router-type
Zone2-SI-S(config-vlan-1)# static-mac-address abcd.5200.0b4e ethernet 16 high-
priority router-type
Zone2-SI-S(config-vlan-1)# exit
Zone2-SI-S(config)# write memory
Zone2-SI-S(config)# exit
Zone2-SI-S# reload
```

Commands on Zone3-SI-A in Zone 3

The following commands configure ServerIron “Zone3-SI-A”, on the left side of zone 3 in Figure 6.2 on page 6-8.

```
ServerIron(config)# hostname Zone3-SI-A
Zone3-SI-A(config)# ip address 209.157.23.11 255.255.255.0
Zone3-SI-A(config)# ip default-gateway 209.157.23.1
Zone3-SI-A(config)# no span
Zone3-SI-A(config)# ip policy 1 fw tcp 0 global
Zone3-SI-A(config)# ip policy 2 fw udp 0 global
Zone3-SI-A(config)# server router-ports 5
Zone3-SI-A(config)# server fw-port 9
Zone3-SI-A(config)# trunk switch ethernet 9 to 10
Zone3-SI-A(config)# vlan 10 by port
Zone3-SI-A(config-vlan-10)# untagged 9 to 10
Zone3-SI-A(config-vlan-10)# exit
Zone3-SI-A(config)# vlan 1
Zone3-SI-A(config-vlan-1)# always-active
Zone3-SI-A(config-vlan-1)# exit
Zone3-SI-A(config)# server fw-name FW1 209.157.23.1
Zone3-SI-A(config-rs-FW1)# exit
Zone3-SI-A(config)# server fw-name FW2 209.157.23.254
Zone3-SI-A(config-rs-FW2)# exit
Zone3-SI-A(config)# access-list 2 permit 209.157.25.0 0.0.0.255
Zone3-SI-A(config)# server fw-group 2
Zone3-SI-A(config-tc-2)# fwall-zone Zone2 2 2
```



```
Zone3-SI-A(config-tc-2)# fw-name FW1
Zone3-SI-A(config-tc-2)# fw-name FW2
Zone3-SI-A(config-tc-2)# l2-fwall
Zone3-SI-A(config-tc-2)# sym-priority 1
Zone3-SI-A(config-tc-2)# fwall-info 1 1 209.157.24.13 209.157.23.1
Zone3-SI-A(config-tc-2)# fwall-info 2 1 209.157.24.14 209.157.23.1
Zone3-SI-A(config-tc-2)# fwall-info 3 16 209.157.24.13 209.157.23.254
Zone3-SI-A(config-tc-2)# fwall-info 4 16 209.157.24.14 209.157.23.254
Zone3-SI-A(config-tc-2)# fwall-info 5 1 209.157.25.15 209.157.23.1
Zone3-SI-A(config-tc-2)# fwall-info 6 1 209.157.25.16 209.157.23.1
Zone3-SI-A(config-tc-2)# fwall-info 7 16 209.157.25.15 209.157.23.254
Zone3-SI-A(config-tc-2)# fwall-info 8 16 209.157.25.16 209.157.23.254
Zone3-SI-A(config-tc-2)# fwall-info 9 5 209.157.23.15 209.157.23.15
Zone3-SI-A(config-tc-2)# exit
Zone3-SI-A(config)# vlan 1
Zone3-SI-A(config-vlan-1)# static-mac-address abcd.5200.3489 ethernet 1 high-
priority router-type
Zone3-SI-A(config-vlan-1)# static-mac-address abcd.5200.0b4c ethernet 16 high-
priority router-type
Zone3-SI-A(config-vlan-1)# exit
Zone3-SI-A(config)# write memory
Zone3-SI-A(config)# exit
Zone3-SI-A# reload
```

Commands on Zone3-SI-S in Zone 3

The following commands configure ServerIron “Zone3-SI-S”, on the right side of zone 3 in Figure 6.2 on page 6-8.

```
ServerIron(config)# hostname Zone3-SI-S
Zone3-SI-S(config)# ip address 209.157.23.12 255.255.255.0
Zone3-SI-S(config)# ip default-gateway 209.157.23.254
Zone3-SI-S(config)# no span
Zone3-SI-S(config)# ip policy 1 fw tcp 0 global
Zone3-SI-S(config)# ip policy 2 fw udp 0 global
Zone3-SI-S(config)# server router-ports 5
Zone3-SI-S(config)# server fw-port 9
Zone3-SI-S(config)# trunk switch ethernet 9 to 10
Zone3-SI-S(config)# vlan 10 by port
Zone3-SI-S(config-vlan-10)# untagged 9 to 10
Zone3-SI-S(config-vlan-10)# exit
Zone3-SI-S(config)# vlan 1
Zone3-SI-S(config-vlan-1)# always-active
Zone3-SI-S(config-vlan-1)# exit
Zone3-SI-S(config)# server fw-name FW1 209.157.23.1
Zone3-SI-S(config-rs-FW1)# exit
Zone3-SI-S(config)# server fw-name FW2 209.157.23.254
Zone3-SI-S(config-rs-FW2)# exit
Zone3-SI-S(config)# access-list 2 permit 209.157.25.0 0.0.0.255
Zone3-SI-S(config)# server fw-group 2
Zone3-SI-S(config-tc-2)# fwall-zone Zone2 2 2
Zone3-SI-S(config-tc-2)# fw-name FW1
Zone3-SI-S(config-tc-2)# fw-name FW2
Zone3-SI-S(config-tc-2)# l2-fwall
Zone3-SI-S(config-tc-2)# sym-priority 1
Zone3-SI-S(config-tc-2)# fwall-info 1 16 209.157.24.13 209.157.23.1
Zone3-SI-S(config-tc-2)# fwall-info 2 16 209.157.24.14 209.157.23.1
Zone3-SI-S(config-tc-2)# fwall-info 3 1 209.157.24.13 209.157.23.254
Zone3-SI-S(config-tc-2)# fwall-info 4 1 209.157.24.14 209.157.23.254
Zone3-SI-S(config-tc-2)# fwall-info 5 16 209.157.25.15 209.157.23.1
```



```
Zone3-SI-S(config-tc-2)# fwall-info 6 16 209.157.25.16 209.157.23.1
Zone3-SI-S(config-tc-2)# fwall-info 7 1 209.157.25.15 209.157.23.1
Zone3-SI-S(config-tc-2)# fwall-info 8 1 209.157.25.16 209.157.23.254
Zone3-SI-S(config-tc-2)# fwall-info 9 5 209.157.23.15 209.157.23.15
Zone3-SI-S(config-tc-2)# exit
Zone3-SI-S(config)# vlan 1
Zone3-SI-S(config-vlan-1)# static-mac-address abcd.5200.3489 ethernet 1
high-priority router-type
Zone3-SI-S(config-vlan-1)# static-mac-address abcd.5200.0b4c ethernet 16
high-priority router-type
Zone3-SI-S(config-vlan-1)# exit
Zone3-SI-S(config)# write memory
Zone3-SI-S(config)# exit
Zone3-SI-S# reload
```

IronClad FWLB configurations require each ServerIron in an active-standby pair to have a link to each of the firewalls for which the ServerIrons are providing load balancing.

If the firewalls are multi-homed (allow more than one connection on each side of the protected network), then it is possible to connect each ServerIron to all the firewalls directly. Figure 6.3 on page 6-18 shows an example of this type of configuration.

Figure 6.3 IronClad FWLB configuration with multi-homed firewalls

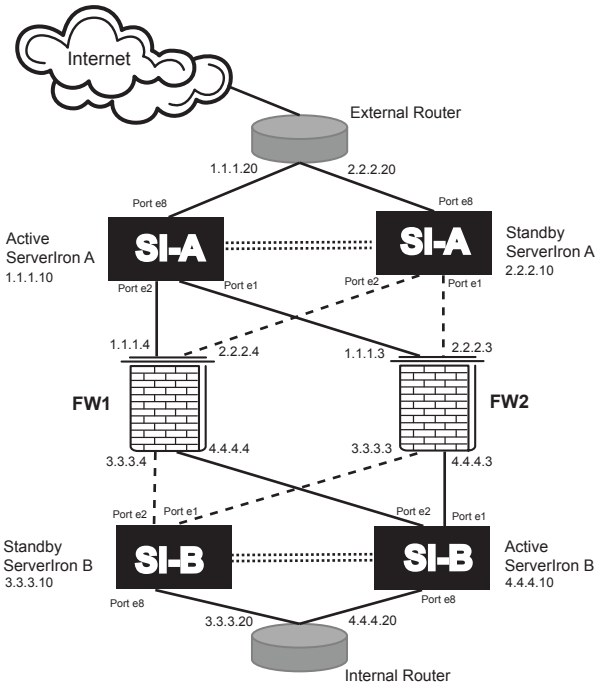

In this example, each firewall has four interfaces. Each interface goes to a ServerIron.

NOTE: If the firewalls are not multi-homed, you need to use additional devices, typically Layer 2 switches, to provide the redundant links. shows an example of an IronClad FWLB configuration that uses Layer 2 switches to provide multi-homing between the ServerIron and firewalls.

Configuration Examples with Layer 3 Routing

NOTE: Layer 3 routing is supported only on ServerIron Chassis devices running software release 08.0.00 or later.

This section shows examples of commonly used ServerIron multizone FWLB deployments with Layer 3 configurations. The ServerIrons in these examples perform Layer 3 routing in addition to Layer 2 and Layer 4 – 7 switching.

Generally, the steps for configuring Layer 4 – 7 features on a ServerIron running Layer 3 are similar to the steps on a ServerIron that is not running Layer 3. The examples focus on the Layer 3 aspects of the configurations.

This section contains the following configuration examples:

- "Multizone FWLB with One Sub-net and One Virtual Routing Interface" on page 6-19
- "Multizone FWLB with Multiple Sub-nets and Multiple Virtual Routing Interfaces" on page 6-28

NOTE: The multizone FWLB configurations shown in these examples are the ones that are supported. If you need to use the ServerIron's Layer 3 routing support in a FWLB configuration that is not shown, contact Brocade Communications Systems.

Multizone FWLB with One Sub-net and One Virtual Routing Interface

Multizone FWLB allows you to configure ServerIrons to forward packets based on the destination zone. For example, if your network consists of an Internet side, an internal side, and a Demilitarized Zone (DMZ) in between, you can configure ServerIrons to forward packets through the firewalls to the correct zone.

When you configure multi-zone FWLB, you first identify a zone by configuring standard ACLs. An ACL specifies the IP addresses (or address ranges) within the zone. When you configure the firewall group parameters, you add the zones and define them by associating the ACLs with them. Each zone consists of a zone number, an optional name, and a standard IP ACL that specifies the IP addresses contained in the zone.

Figure 6.4 shows an example of a multizone configuration for three zones:

- Zone 1 – The default zone. All sub-nets that you do not configure to be members of the other zones are by default members of zone 1. Generally, the default zone is on the public (non-secure) side of the firewalls.
- Zone 2 – A secured zone containing two application servers.
- Zone 3 – Another secured zone containing an additional application server.

The ServerIrons in zone 1 perform FWLB for traffic between zone 1 and zones 2 and 3.

Figure 6.4 Multizone FWLB with One Sub-net and One Virtual Routing Interface

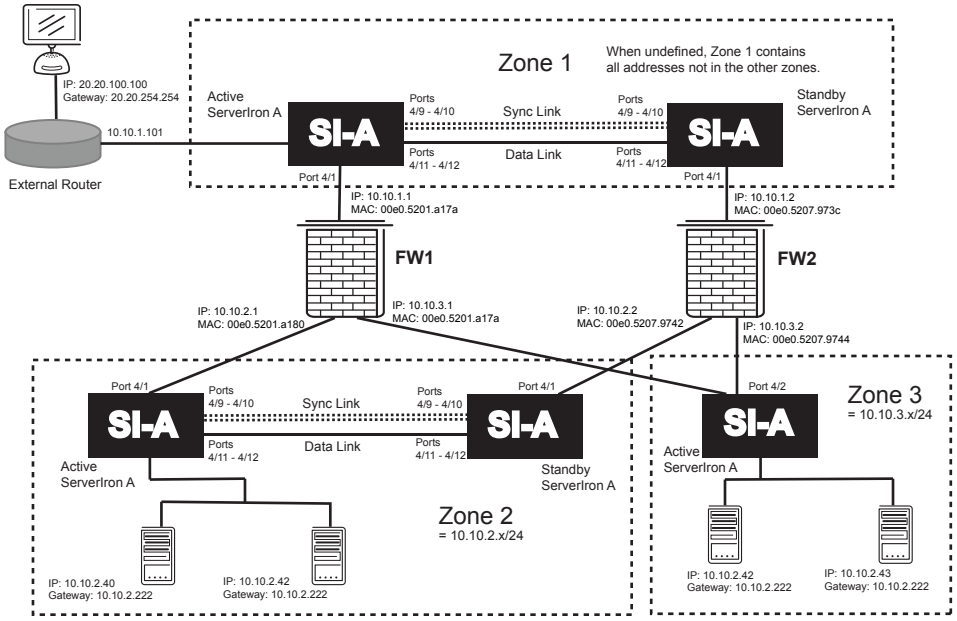

This configuration example also uses SLB. The application servers connected to the ServerIrons in zones 2 and 3 are configured on the ServerIrons as real servers and bound to a VIP. The ServerIrons in zone 1 load balance client requests for the servers in zones 2 and 3, in addition to load balancing the traffic to the firewalls. FWLB-to-SLB and SLB-to-FWLB are used in this configuration. FWLB-to-SLB enables the ServerIrons in zones 2 and 3 to learn the firewall from which a client request is received and send the server reply back through the same firewall. SLB-to-FWLB on the ServerIrons in zone 1 performs FWLB for traffic directed toward the real servers connected to the ServerIrons in zones 2 and 3.

Commands on Zone 1's Active ServerIron (Zone1-SI-A)

The following commands change the CLI to the global CONFIG level, then change the hostname to "Zone1-SI-A".

```
ServerIron> enable
ServerIron# configure terminal
ServerIron(config)# hostname Zone1-SI-A
```

The following commands enable the always-active feature and disable the Spanning Tree Protocol (STP) in VLAN 1, which contains the ports that will carry the FWLB traffic.

```
Zone1-SI-A(config)# vlan 1
Zone1-SI-A(config-vlan-1)# always-active
Zone1-SI-A(config-vlan-1)# no spanning-tree
```

The following commands configure a virtual routing interface on VLAN 1 (the default VLAN), then configure an IP address on the interface. The virtual routing interface is associated with all the ports in the VLAN.

```
Zone1-SI-A(config-vlan-1)# router-interface ve 1
Zone1-SI-A(config-vlan-1)# exit
Zone1-SI-A(config)# interface ve 1
Zone1-SI-A(config-ve-1)# ip address 10.10.1.111 255.255.255.0
Zone1-SI-A(config-ve-1)# exit
```

The following command configures an IP default route. The next hop for this route is the ServerIron's interface with firewall FW1.

```
Zone1-SI-A(config)# ip route 0.0.0.0 0.0.0.0 10.10.1.1
```

The following command disables ICMP redirect messages. This command disables the messages but the ServerIron still forwards misdirected traffic to the appropriate router.

```
Zone1-SI-A(config)# no ip icmp redirects
```


The following commands configure the synchronization link between this ServerIron and ServerIron Zone1-SI-B. For redundancy, the link is configured on a trunk group.

```
Zone1-SI-A(config)# vlan 10
Zone1-SI-A(config-vlan-10)# untagged ethernet 4/9 to 4/10
Zone1-SI-A(config-vlan-10)# exit
Zone1-SI-A(config)# trunk switch ethernet 4/9 to 4/10
Zone1-SI-A(config)# server fw-port 4/9
```

The following commands configure the data link connecting this ServerIron to its partner, Zone1-SI-B. For redundancy, the link is configured as a two-port trunk group.

```
Zone1-SI-A(config)# trunk switch ethernet 4/11 to 4/12
Zone1-SI-A(config)# server partner-ports ethernet 4/11
Zone1-SI-A(config)# server partner-ports ethernet 4/12
Zone1-SI-A(config)# server fw-group 2
Zone1-SI-A(config-tc-2)# l2-fwall
Zone1-SI-A(config-tc-2)# exit
```

The following commands add the firewalls. Three application ports (HTTP, FTP, and SNMP) are configured on each of the firewalls. The no-health-check parameter disables the Layer 4 health check for the specified application.

```
Zone1-SI-A(config)# server fw-name fw1 10.10.1.1
Zone1-SI-A(config-rs-fw1)# port http
Zone1-SI-A(config-rs-fw1)# port http no-health-check
Zone1-SI-A(config-rs-fw1)# port ftp
Zone1-SI-A(config-rs-fw1)# port ftp no-health-check
Zone1-SI-A(config-rs-fw1)# port snmp
Zone1-SI-A(config-rs-fw1)# port snmp no-health-check
Zone1-SI-A(config-rs-fw1)# exit
Zone1-SI-A(config)# server fw-name fw2 10.10.1.2
Zone1-SI-A(config-rs-fw2)# port http
Zone1-SI-A(config-rs-fw2)# port http no-health-check
Zone1-SI-A(config-rs-fw2)# port ftp
Zone1-SI-A(config-rs-fw2)# port ftp no-health-check
Zone1-SI-A(config-rs-fw2)# port snmp
Zone1-SI-A(config-rs-fw2)# port snmp no-health-check
Zone1-SI-A(config-rs-fw2)# exit
```

The following commands add the firewall definitions to the firewall port group (always group 2). The firewall group contains all the ports in VLAN 1 (the default VLAN).

```
Zone1-SI-A(config)# server fw-group 2
Zone1-SI-A(config-tc-2)# fw-name fw1
Zone1-SI-A(config-tc-2)# fw-name fw2
```

The following command enables the active-active mode and specifies the priority of this ServerIron. In this case, ServerIron Zone1-SI-A has the higher priority. Its partner, ServerIron Zone1-SI-B, will be configured with a lower priority (1).

```
Zone1-SI-A(config-tc-2)# sym-priority 255
```

The following commands add the paths through the firewalls to the ServerIrons in zones 2 and 3. In addition, static MAC entries are added for the firewall interfaces. Static MAC entries are required in this type of configuration, in which one sub-net and one virtual routing interface are used.

NOTE: The path IDs must be in contiguous, ascending numerical order, starting with 1. For example, path sequence 1, 2, 3, 4 is valid. Path sequence 4, 3, 2, 1 or 1, 3, 4, 5 is not valid.

```
Zone1-SI-A(config-tc-2)# fwall-info 1 4/1 10.10.2.222 10.10.1.1
Zone1-SI-A(config-tc-2)# fwall-info 2 4/11 10.10.2.222 10.10.1.2
Zone1-SI-A(config-tc-2)# fwall-info 3 4/1 10.10.2.223 10.10.1.1
Zone1-SI-A(config-tc-2)# fwall-info 4 4/11 10.10.2.223 10.10.1.2
Zone1-SI-A(config-tc-2)# fwall-info 5 4/1 10.10.3.111 10.10.1.1
```



```
Zone1-SI-A(config-tc-2)# fwall-info 6 4/11 10.10.3.111 10.10.1.2
Zone1-SI-A(config-tc-2)# exit
Zone1-SI-A(config)# vlan 1
Zone1-SI-A(config-vlan-1)# static-mac-address 00e0.5201.a17a ethernet 4/1 priority 1
router-type
Zone1-SI-A(config-vlan-1)# static-mac-address 00e0.5207.973c ethernet 4/11 priority
1 router-type
Zone1-SI-A(config-vlan-1)# exit
```

The following commands set the load balancing method to balance requests based on the firewall that has the least number of connections for the requested service. For example, the ServerIron will load balance HTTP requests based on the firewall that has fewer HTTP session entries in the ServerIron session table.

```
Zone1-SI-A(config)# server fw-group 2
Zone1-SI-A(config-tc-2)# fw-predictor per-service-least-conn
Zone1-SI-A(config-tc-2)# exit
```

The following command configures a standard IP ACL for the IP addresses in one of the zones this ServerIron is not in. In this configuration, only one zone definition is required on each ServerIron, including Zone1-SI-A and Zone1-SI-S. Since the active Zone 1 ServerIron is already in zone 1, the ServerIron will forward packets either to the active ServerIron in zone 2 or to the only other active ServerIron that is not in zone 2. In this case, the other active ServerIron is in zone 3. Thus, if ServerIron Zone1-SI-A receives a packet that is not addressed to the sub-net Zone1-SI-A is in, and is not addressed to a sub-net in zone 2, the ServerIron assumes that the packet is for an address in the other zone, zone 3. The ServerIron forwards the packet to the ServerIron in zone 3.

The command configures an ACL for the addresses in zone 2, which contains addresses in the 10.10.2.x/24 sub-net. The “0.0.0.255” values indicate the significant bits in the IP address you specify. In this case, all bits except the ones in the last node of the address are significant.

```
Zone1-SI-A(config)# access-list 2 permit 10.10.2.0 0.0.0.255
```

The following commands configure the zone parameters. To configure a zone, specify a name for the zone, then a zone number (from 1 – 10), followed by the number of the ACL that specifies the IP addresses in the zone. In this example, the ACL numbers and zone numbers are the same, but this is not required.

```
Zone1-SI-A(config)# server fw-group 2
Zone1-SI-A(config-tc-2)# fwall-zone Zone2 2 2
Zone1-SI-A(config-tc-2)# exit
```

The following commands configure the SLB information. Each of the servers in zones 2 and 3 is added as a real server, then the servers are bound to a VIP. The servers are added using the **server remote-name** command instead of the **server real-name** command because the servers are not directly connected to the ServerIron. Instead, they are connected to the ServerIron through other routers (in this case, the firewalls).

```
Zone1-SI-A(config)# server remote-name web1 10.10.2.40
Zone1-SI-A(config-rs-web1)# port http
Zone1-SI-A(config-rs-web1)# exit
Zone1-SI-A(config)# server remote-name web2 10.10.2.42
Zone1-SI-A(config-rs-web2)# port http
Zone1-SI-A(config-rs-web2)# exit
Zone1-SI-A(config)# server remote-name web3 10.10.3.41
Zone1-SI-A(config-rs-web3)# port http
Zone1-SI-A(config-rs-web3)# exit
Zone1-SI-A(config)# server remote-name web4 10.10.3.43
Zone1-SI-A(config-rs-web4)# port http
Zone1-SI-A(config-rs-web4)# exit
Zone1-SI-A(config)# server virtual www.web.com 10.10.1.10
Zone1-SI-A(config-vs-www.web.com)# port http
Zone1-SI-A(config-vs-www.web.com)# bind http web1 http web2 http web3 http web4 http
Zone1-SI-A(config-vs-www.web.com)# exit
```

The following command enables SLB-to-FWLB.

```
Zone1-SI-A(config)# server slb-fw
```

The following commands enable FWLB.


```
Zone1-SI-A(config)# ip l4-policy 1 fw tcp 0 global
Zone1-SI-A(config)# ip l4-policy 2 fw udp 0 global
```

The following command saves the configuration changes to the startup-config file.

```
Zone1-SI-A(config)# write memory
```

Commands on Zone 1's Standby ServerIron (Zone1-SI-S)

```
ServerIron> enable
ServerIron# configure terminal
ServerIron(config)# hostname Zone1-SI-S
Zone1-SI-S(config)# vlan 1
Zone1-SI-S(config-vlan-1)# always-active
Zone1-SI-S(config-vlan-1)# no spanning-tree
Zone1-SI-S(config-vlan-1)# router-interface ve 1
Zone1-SI-S(config-vlan-1)# exit
Zone1-SI-S(config)# interface ve 1
Zone1-SI-S(config-ve-1)# ip address 10.10.1.112 255.255.255.0
Zone1-SI-S(config-ve-1)# exit
Zone1-SI-S(config)# ip route 0.0.0.0 0.0.0.0 10.10.1.2
Zone1-SI-S(config)# no ip icmp redirects
Zone1-SI-S(config)# vlan 10
Zone1-SI-S(config-vlan-10)# untagged ethernet 4/9 to 4/10
Zone1-SI-S(config-vlan-10)# exit
Zone1-SI-S(config)# trunk switch ethernet 4/9 to 4/10
Zone1-SI-S(config)# server fw-port 4/9
Zone1-SI-S(config)# trunk switch ethernet 4/11 to 4/12
Zone1-SI-S(config)# server partner-ports ethernet 4/11
Zone1-SI-S(config)# server partner-ports ethernet 4/12
Zone1-SI-S(config)# server fw-group 2
Zone1-SI-S(config-tc-2)# l2-fwall
Zone1-SI-S(config-tc-2)# exit
Zone1-SI-S(config)# server fw-name fw1 10.10.1.1
Zone1-SI-S(config-rs-fw1)# port http
Zone1-SI-S(config-rs-fw1)# port http no-health-check
Zone1-SI-S(config-rs-fw1)# port ftp
Zone1-SI-S(config-rs-fw1)# port ftp no-health-check
Zone1-SI-S(config-rs-fw1)# port snmp
Zone1-SI-S(config-rs-fw1)# port snmp no-health-check
Zone1-SI-S(config-rs-fw1)# exit
Zone1-SI-S(config)# server fw-name fw2 10.10.1.2
Zone1-SI-S(config-rs-fw2)# port http
Zone1-SI-S(config-rs-fw2)# port http no-health-check
Zone1-SI-S(config-rs-fw2)# port ftp
Zone1-SI-S(config-rs-fw2)# port ftp no-health-check
Zone1-SI-S(config-rs-fw2)# port snmp
Zone1-SI-S(config-rs-fw2)# port snmp no-health-check
Zone1-SI-S(config-rs-fw2)# exit
Zone1-SI-S(config)# server fw-group 2
Zone1-SI-S(config-tc-2)# fw-name fw1
Zone1-SI-S(config-tc-2)# fw-name fw2
Zone1-SI-S(config-tc-2)# sym-priority 1
Zone1-SI-S(config-tc-2)# fwall-info 1 4/11 10.10.2.222 10.10.1.1
Zone1-SI-S(config-tc-2)# fwall-info 2 4/1 10.10.2.222 10.10.1.2
Zone1-SI-S(config-tc-2)# fwall-info 3 4/11 10.10.2.223 10.10.1.1
Zone1-SI-S(config-tc-2)# fwall-info 4 4/1 10.10.2.223 10.10.1.2
Zone1-SI-S(config-tc-2)# fwall-info 5 4/11 10.10.3.111 10.10.1.1
Zone1-SI-S(config-tc-2)# fwall-info 6 4/1 10.10.3.111 10.10.1.2
Zone1-SI-S(config-tc-2)# exit
```



```
Zone1-SI-S(config)# vlan 1
Zone1-SI-S(config-vlan-1)# static-mac-address 00e0.5201.a17a ethernet 4/11 priority
1 router-type
Zone1-SI-S(config-vlan-1)# static-mac-address 00e0.5207.973c ethernet 4/1 priority 1
router-type
Zone1-SI-S(config-vlan-1)# exit
Zone1-SI-S(config-tc-2)# server fw-group 2
Zone1-SI-S(config-tc-2)# fw-predictor per-service-least-conn
Zone1-SI-S(config-tc-2)# exit
Zone1-SI-S(config)# access-list 2 permit 10.10.2.0 0.0.0.255
Zone1-SI-S(config)# server fw-group 2
Zone1-SI-S(config-tc-2)# fwall-zone Zone2 2 2
Zone1-SI-S(config-tc-2)# exit
Zone1-SI-S(config)# server remote-name web1 10.10.2.40
Zone1-SI-S(config-rs-web1)# port http
Zone1-SI-S(config-rs-web1)# exit
Zone1-SI-S(config)# server remote-name web2 10.10.2.42
Zone1-SI-S(config-rs-web2)# port http
Zone1-SI-S(config-rs-web2)# exit
Zone1-SI-S(config)# server remote-name web3 10.10.3.41
Zone1-SI-S(config-rs-web3)# port http
Zone1-SI-S(config-rs-web3)# exit
Zone1-SI-S(config)# server remote-name web4 10.10.3.43
Zone1-SI-S(config-rs-web4)# port http
Zone1-SI-S(config-rs-web4)# exit
Zone1-SI-S(config)# server virtual www.web.com 10.10.1.10
Zone1-SI-S(config-vs-www.web.com)# port http
Zone1-SI-S(config-vs-www.web.com)# bind http web1 http web2 http web3 http web4 http
Zone1-SI-S(config-vs-www.web.com)# exit
Zone1-SI-S(config)# server slb-fw
Zone1-SI-S(config)# ip l4-policy 1 fw tcp 0 global
Zone1-SI-S(config)# ip l4-policy 2 fw udp 0 global
Zone1-SI-S(config)# write memory
```

Commands on Zone 2's Active ServerIron (Zone2-SI-A)

The following commands configure ServerIron Zone2-SI-A in zone 2. The configuration is similar to the configuration for ServerIron Zone1-SI-A, except the ACL and zone information are for zone 3, and FWLB-to-SLB is enabled instead of SLB-to-FWLB.

```
ServerIron> enable
ServerIron# configure terminal
ServerIron(config)# hostname Zone2-SI-A
Zone2-SI-A(config)# vlan 1
Zone2-SI-A(config-vlan-1)# always-active
Zone2-SI-A(config-vlan-1)# no spanning-tree
Zone2-SI-A(config-vlan-1)# router-interface ve 1
Zone2-SI-A(config-vlan-1)# exit
Zone2-SI-A(config)# interface ve 1
Zone2-SI-A(config-ve-1)# ip address 10.10.2.222 255.255.255.0
Zone2-SI-A(config-ve-1)# exit
Zone2-SI-A(config)# ip route 0.0.0.0 0.0.0.0 10.10.2.1
Zone2-SI-A(config)# no ip icmp redirects
Zone2-SI-A(config)# vlan 10
Zone2-SI-A(config-vlan-10)# untagged ethernet 4/9 to 4/10
Zone2-SI-A(config-vlan-10)# exit
Zone2-SI-A(config)# trunk switch ethernet 4/9 to 4/10
Zone2-SI-A(config)# server fw-port 4/9
Zone2-SI-A(config)# trunk switch ethernet 4/11 to 4/12
```



```
Zone2-SI-A(config)# server partner-ports ethernet 4/11
Zone2-SI-A(config)# server partner-ports ethernet 4/12
Zone2-SI-A(config)# server fw-group 2
Zone2-SI-A(config-tc-2)# l2-fwall
Zone2-SI-A(config-tc-2)# exit
Zone2-SI-A(config)# server fw-name fw1 10.10.2.1
Zone2-SI-A(config-rs-fw1)# port http
Zone2-SI-A(config-rs-fw1)# port http no-health-check
Zone2-SI-A(config-rs-fw1)# port ftp
Zone2-SI-A(config-rs-fw1)# port ftp no-health-check
Zone2-SI-A(config-rs-fw1)# port snmp
Zone2-SI-A(config-rs-fw1)# port snmp no-health-check
Zone2-SI-A(config-rs-fw1)# exit
Zone2-SI-A(config)# server fw-name fw2 10.10.2.2
Zone2-SI-A(config-rs-fw2)# port http
Zone2-SI-A(config-rs-fw2)# port http no-health-check
Zone2-SI-A(config-rs-fw2)# port ftp
Zone2-SI-A(config-rs-fw2)# port ftp no-health-check
Zone2-SI-A(config-rs-fw2)# port snmp
Zone2-SI-A(config-rs-fw2)# port snmp no-health-check
Zone2-SI-A(config-rs-fw2)# exit
Zone2-SI-A(config)# server fw-group 2
Zone2-SI-A(config-tc-2)# fw-name fw1
Zone2-SI-A(config-tc-2)# fw-name fw2
Zone2-SI-A(config-tc-2)# sym-priority 255
Zone2-SI-A(config-tc-2)# fwall-info 1 4/1 10.10.1.111 10.10.2.1
Zone2-SI-A(config-tc-2)# fwall-info 2 4/11 10.10.1.111 10.10.2.2
Zone2-SI-A(config-tc-2)# fwall-info 3 4/1 10.10.1.112 10.10.2.1
Zone2-SI-A(config-tc-2)# fwall-info 4 4/11 10.10.1.112 10.10.2.2
Zone2-SI-A(config-tc-2)# fwall-info 5 4/1 10.10.3.111 10.10.2.1
Zone2-SI-A(config-tc-2)# fwall-info 6 4/11 10.10.3.111 10.10.2.2
Zone2-SI-A(config-tc-2)# exit
Zone2-SI-A(config)# vlan 1
Zone2-SI-A(config-vlan-1)# static-mac-address 00e0.5201.a180 ethernet 4/1 priority 1
router-type
Zone2-SI-A(config-vlan-1)# static-mac-address 00e0.5207.9742 ethernet 4/11 priority
1 router-type
Zone2-SI-A(config-vlan-1)# exit
Zone2-SI-A(config)# server fw-group 2
Zone2-SI-A(config-tc-2)# fw-predictor per-service-least-conn
Zone2-SI-A(config-tc-2)# exit
Zone2-SI-A(config)# access-list 3 permit 10.10.3.0 0.0.0.255
Zone2-SI-A(config)# server fw-group 2
Zone2-SI-A(config-tc-2)# fwall-zone zone3 3 3
Zone2-SI-A(config-tc-2)# exit
Zone2-SI-A(config)# server real-name rs1 10.10.2.40
Zone2-SI-A(config-rs-rs1)# port http
Zone2-SI-A(config-rs-rs1)# exit
Zone2-SI-A(config)# server real-name rs1 10.10.2.42
Zone2-SI-A(config-rs-rs2)# port http
Zone2-SI-A(config-rs-rs2)# exit
Zone2-SI-A(config)# server virtual www.rs.com 10.10.2.10
Zone2-SI-A(config-vs-www.rs.com)# port http
Zone2-SI-A(config-vs-www.web.com)# bind http rs1 http rs2 http
Zone2-SI-A(config-vs-www.web.com)# exit
Zone2-SI-A(config)# server fw-slb
Zone2-SI-A(config)# ip l4-policy 1 fw tcp 0 global
Zone2-SI-A(config)# ip l4-policy 2 fw udp 0 global
Zone2-SI-A(config)# write memory
```


Commands on Zone 2's Standby ServerIron (Zone2-SI-S)

```
ServerIron> enable
ServerIron# configure terminal
ServerIron(config)# hostname Zone2-SI-S
Zone2-SI-S(config)# vlan 1
Zone2-SI-S(config-vlan-1)# always-active
Zone2-SI-S(config-vlan-1)# no spanning-tree
Zone2-SI-S(config-vlan-1)# router-interface ve 1
Zone2-SI-S(config-vlan-1)# exit
Zone2-SI-S(config)# interface ve 1
Zone2-SI-S(config-ve-1)# ip address 10.10.2.223 255.255.255.0
Zone2-SI-S(config-ve-1)# exit
Zone2-SI-S(config)# ip route 0.0.0.0 0.0.0.0 10.10.2.2
Zone2-SI-S(config)# no ip icmp redirects
Zone2-SI-S(config)# vlan 10
Zone2-SI-S(config-vlan-10)# untagged ethernet 4/9 to 4/10
Zone2-SI-S(config-vlan-10)# exit
Zone2-SI-S(config)# trunk switch ethernet 4/9 to 4/10
Zone2-SI-S(config)# server fw-port 4/9
Zone2-SI-S(config)# trunk switch ethernet 4/11 to 4/12
Zone2-SI-S(config)# server partner-ports ethernet 4/11
Zone2-SI-S(config)# server partner-ports ethernet 4/12
Zone2-SI-S(config)# server fw-group 2
Zone2-SI-S(config-tc-2)# l2-fwall
Zone2-SI-S(config-tc-2)# exit
Zone2-SI-S(config)# server fw-name fw1 10.10.2.1
Zone2-SI-S(config-rs-fw1)# port http
Zone2-SI-S(config-rs-fw1)# port http no-health-check
Zone2-SI-S(config-rs-fw1)# port ftp
Zone2-SI-S(config-rs-fw1)# port ftp no-health-check
Zone2-SI-S(config-rs-fw1)# port snmp
Zone2-SI-S(config-rs-fw1)# port snmp no-health-check
Zone2-SI-S(config-rs-fw1)# exit
Zone2-SI-S(config)# server fw-name fw2 10.10.2.2
Zone2-SI-S(config-rs-fw2)# port http
Zone2-SI-S(config-rs-fw2)# port http no-health-check
Zone2-SI-S(config-rs-fw2)# port ftp
Zone2-SI-S(config-rs-fw2)# port ftp no-health-check
Zone2-SI-S(config-rs-fw2)# port snmp
Zone2-SI-S(config-rs-fw2)# port snmp no-health-check
Zone2-SI-S(config-rs-fw2)# exit
Zone2-SI-S(config)# server fw-group 2
Zone2-SI-S(config-tc-2)# fw-name fw1
Zone2-SI-S(config-tc-2)# fw-name fw2
Zone2-SI-S(config-tc-2)# sym-priority 1
Zone2-SI-S(config-tc-2)# fwall-info 1 4/11 10.10.1.111 10.10.2.1
Zone2-SI-S(config-tc-2)# fwall-info 2 4/1 10.10.1.111 10.10.2.2
Zone2-SI-S(config-tc-2)# fwall-info 3 4/11 10.10.1.112 10.10.2.1
Zone2-SI-S(config-tc-2)# fwall-info 4 4/1 10.10.1.112 10.10.2.2
Zone2-SI-S(config-tc-2)# fwall-info 5 4/11 10.10.3.111 10.10.2.1
Zone2-SI-S(config-tc-2)# fwall-info 6 4/1 10.10.3.111 10.10.2.2
Zone2-SI-S(config-tc-2)# exit
Zone2-SI-S(config)# vlan 1
Zone2-SI-S(config-vlan-1)# static-mac-address 00e0.5201.a180 ethernet 4/11 priority
1 router-type
Zone2-SI-S(config-vlan-1)# static-mac-address 00e0.5207.9742 ethernet 4/1 priority 1
router-type
Zone2-SI-S(config-vlan-1)# exit
```



```
Zone2-SI-S(config)# server group 2
Zone2-SI-S(config-tc-2)# fw-predictor per-service-least-conn
Zone2-SI-S(config-tc-2)# exit
Zone2-SI-S(config)# access-list 3 permit 10.10.3.0 0.0.0.255
Zone2-SI-S(config)# server fw-group 2
Zone2-SI-S(config-tc-2)# fwall-zone zone3 3 3
Zone2-SI-S(config-tc-2)# exit
Zone2-SI-S(config)# server real-name rs1 10.10.2.40
Zone2-SI-S(config-rs-rs1)# port http
Zone2-SI-S(config-rs-rs1)# exit
Zone2-SI-S(config)# server real-name rs1 10.10.2.42
Zone2-SI-S(config-rs-rs2)# port http
Zone2-SI-S(config-rs-rs2)# exit
Zone2-SI-S(config)# server virtual www.rs.com 10.10.2.10
Zone2-SI-S(config-vs-www.rs.com)# port http
Zone2-SI-S(config-vs-www.web.com)# bind http rs1 http rs2 http
Zone2-SI-S(config-vs-www.web.com)# exit
Zone2-SI-S(config)# server fw-slb
Zone2-SI-S(config)# ip l4-policy 1 fw tcp 0 global
Zone2-SI-S(config)# ip l4-policy 2 fw udp 0 global
Zone2-SI-S(config)# write memory
```

Commands on Zone 3's ServerIron (Zone3-SI-A)

Here are the commands for configuring the ServerIron in zone 3.

```
ServerIron> enable
ServerIron# configure terminal
ServerIron(config)# hostname Zone3-SI-A
Zone3-SI-A(config)# vlan 1
Zone3-SI-A(config-vlan-1)# always-active
Zone3-SI-A(config-vlan-1)# no spanning-tree
Zone3-SI-A(config-vlan-1)# router-interface ve 1
Zone3-SI-A(config-vlan-1)# exit
Zone3-SI-A(config)# interface ve 1
Zone3-SI-A(config-ve-1)# ip address 10.10.3.111 255.255.255.0
Zone3-SI-A(config-ve-1)# exit
Zone3-SI-A(config)# ip route 0.0.0.0 0.0.0.0 10.10.3.1
Zone3-SI-A(config)# no ip icmp redirects
Zone3-SI-A(config)# server fw-name fw1 10.10.3.1
Zone3-SI-A(config-rs-fw1)# port http
Zone3-SI-A(config-rs-fw1)# port http no-health-check
Zone3-SI-A(config-rs-fw1)# port ftp
Zone3-SI-A(config-rs-fw1)# port ftp no-health-check
Zone3-SI-A(config-rs-fw1)# port snmp
Zone3-SI-A(config-rs-fw1)# port snmp no-health-check
Zone3-SI-A(config-rs-fw1)# exit
Zone3-SI-A(config)# server fw-name fw2 10.10.3.2
Zone3-SI-A(config-rs-fw2)# port http
Zone3-SI-A(config-rs-fw2)# port http no-health-check
Zone3-SI-A(config-rs-fw2)# port ftp
Zone3-SI-A(config-rs-fw2)# port ftp no-health-check
Zone3-SI-A(config-rs-fw2)# port snmp
Zone3-SI-A(config-rs-fw2)# port snmp no-health-check
Zone3-SI-A(config-rs-fw2)# exit
Zone3-SI-A(config)# server fw-group 2
Zone3-SI-A(config-tc-2)# fw-name fw1
Zone3-SI-A(config-tc-2)# fw-name fw2
Zone3-SI-A(config-tc-2)# fwall-info 1 4/1 10.10.1.111 10.10.3.1
```



```
Zone3-SI-A(config-tc-2)# fwall-info 2 4/2 10.10.1.111 10.10.3.2
Zone3-SI-A(config-tc-2)# fwall-info 3 4/1 10.10.1.112 10.10.3.1
Zone3-SI-A(config-tc-2)# fwall-info 4 4/2 10.10.1.112 10.10.3.2
Zone3-SI-A(config-tc-2)# fwall-info 5 4/1 10.10.2.222 10.10.3.1
Zone3-SI-A(config-tc-2)# fwall-info 6 4/2 10.10.2.222 10.10.3.2
Zone3-SI-A(config-tc-2)# fwall-info 7 4/1 10.10.2.223 10.10.3.1
Zone3-SI-A(config-tc-2)# fwall-info 8 4/2 10.10.2.223 10.10.3.2
Zone3-SI-A(config-tc-2)# exit
Zone3-SI-A(config)# vlan 1
Zone3-SI-A(config-vlan-1)# static-mac-address 00e0.5201.a182 ethernet 4/1 priority 1
router-type
Zone3-SI-A(config-vlan-1)# static-mac-address 00e0.5207.9744 ethernet 4/2 priority 1
router-type
Zone3-SI-A(config-vlan-1)# exit
Zone3-SI-A(config)# server fw-group 2
Zone3-SI-A(config-tc-2)# fw-predictor per-service-least-conn
Zone3-SI-A(config-tc-2)# exit
Zone3-SI-A(config)# access-list 2 permit 10.10.2.0 0.0.0.255
Zone3-SI-A(config)# server fw-group 2
Zone3-SI-A(config-tc-2)# fwall-zone zone2 2 2
Zone3-SI-A(config-tc-2)# exit
Zone3-SI-A(config)# server real-name sr1 10.10.3.41
Zone3-SI-A(config-rs-sr1)# port http
Zone3-SI-A(config-rs-sr1)# exit
Zone3-SI-A(config)# server real-name sr2 10.10.3.43
Zone3-SI-A(config-rs-sr2)# port http
Zone3-SI-A(config-rs-sr2)# exit
Zone3-SI-A(config)# server virtual www.sr.com 10.10.3.10
Zone3-SI-A(config-vs-www.rs.com)# port http
Zone3-SI-A(config-vs-www.web.com)# bind http sr2 http sr1 http
Zone3-SI-A(config-vs-www.web.com)# exit
Zone3-SI-A(config)# server fw-slb
Zone3-SI-A(config)# ip l4-policy 1 fw tcp 0 global
Zone3-SI-A(config)# ip l4-policy 2 fw udp 0 global
Zone3-SI-A(config)# write memory
```

Multizone FWLB with Multiple Sub-nets and Multiple Virtual Routing Interfaces

Figure 6.5 shows an example of a multizone FWLB configuration in which each ServerIron is configured with multiple sub-nets and multiple virtual routing interfaces. The configuration is similar to the one in Figure 6.4 on page 6-20, but differs in the following ways:

- The ServerIrons configured in active-active pairs have four port-based VLANs. VLAN 10 is for the synchronization link between the ServerIrons. The default VLAN (VLAN 1) is not configured with a routing interface. VLANs 2 and 20 are configured with virtual routing interfaces.
- The ServerIrons in zone 1 are configured with a static IP route to the sub-net that the external client is on.
- Static MAC entries are not required and thus are not included for the firewall interfaces.
- More than one standard IP ACL is configured on each ServerIron, since more than one sub-net is a member of each zone.

Figure 6.5 Multizone FWLB with Multiple Sub-nets and Multiple Virtual Routing Interfaces

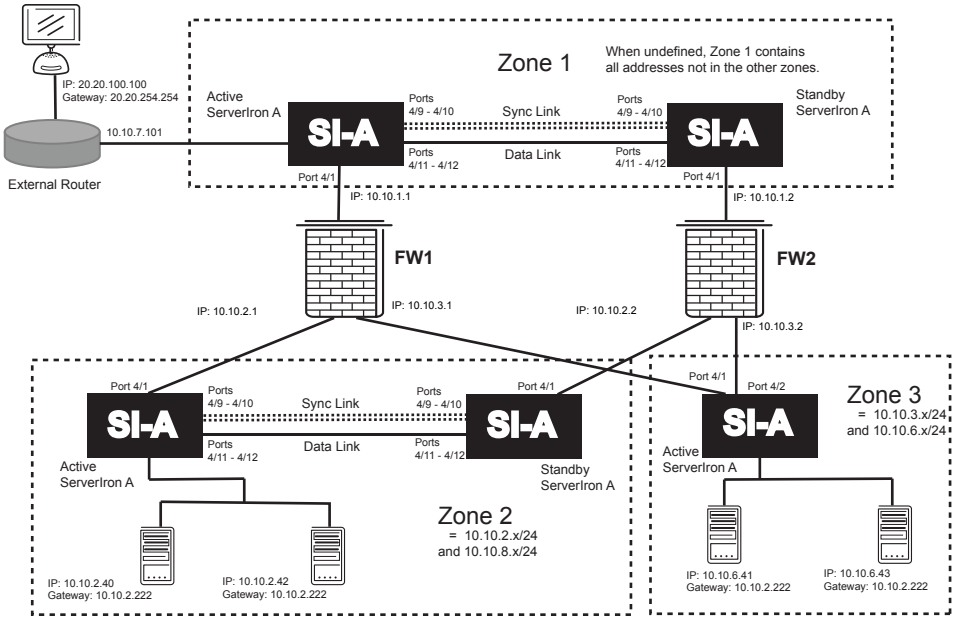

Commands on Zone 1's Active ServerIron (Zone1-SI-A)

```
ServerIron> enable
ServerIron# configure terminal
ServerIron(config)# hostname Zone1-SI-A
```

The following commands enable the always-active feature in VLAN 1.

```
Zone1-SI-A(config)# vlan 1
Zone1-SI-A(config-vlan-1)# always-active
Zone1-SI-A(config-vlan-1)# exit
```

The following commands configure VLAN 2 and virtual routing interface 1, for 10.10.1.111.

```
Zone1-SI-A(config)# vlan 2
Zone1-SI-A(config-vlan-2)# always-active
Zone1-SI-A(config-vlan-2)# tagged ethernet 4/11 to 4/12
Zone1-SI-A(config-vlan-2)# untagged ethernet 4/1 to 4/8
Zone1-SI-A(config-vlan-2)# router-interface ve 1
Zone1-SI-A(config-vlan-2)# exit
Zone1-SI-A(config)# interface ve 1
Zone1-SI-A(config-ve-1)# ip address 10.10.1.111 255.255.255.0
Zone1-SI-A(config-ve-1)# exit
```

The following commands configure VLAN 20 and virtual routing interface 2, for 10.10.7.101.

```
Zone1-SI-A(config)# vlan 20
Zone1-SI-A(config-vlan-20)# always-active
Zone1-SI-A(config-vlan-20)# tagged ethernet 4/11 to 4/12
Zone1-SI-A(config-vlan-20)# untagged ethernet 4/13 to 4/24
Zone1-SI-A(config-vlan-20)# router-interface ve 2
Zone1-SI-A(config-vlan-20# exit
Zone1-SI-A(config)# interface ve 2
Zone1-SI-A(config-ve-2)# ip address 10.10.7.101 255.255.255.0
Zone1-SI-A(config-ve-2)# exit
```

The following command configures an IP default route. The next hop for this route is the ServerIron's interface with firewall FW1.

```
Zone1-SI-A(config)# ip route 0.0.0.0 0.0.0.0 10.10.1.1
```

The following command configures a static route to the sub-net that contains the external host.

```
Zone1-SI-A(config)# ip route 20.20.0.0 255.255.0.0 10.10.7.100
```

The following commands configure the synchronization link between this ServerIron and ServerIron Zone1-SI-B. For redundancy, the link is configured on a trunk group.

```
Zone1-SI-A(config)# vlan 10
Zone1-SI-A(config-vlan-10)# untagged ethernet 4/9 to 4/10
Zone1-SI-A(config-vlan-10)# exit
Zone1-SI-A(config)# trunk switch ethernet 4/9 to 4/10
Zone1-SI-A(config)# server fw-port 4/9
```

The following commands configure the data link connecting this ServerIron to its partner, Zone1-SI-B. For redundancy, the link is configured as a two-port trunk group.

```
Zone1-SI-A(config)# trunk switch ethernet 4/11 to 4/12
Zone1-SI-A(config)# server partner-ports ethernet 4/11
Zone1-SI-A(config)# server partner-ports ethernet 4/12
Zone1-SI-A(config)# server fw-group 2
Zone1-SI-A(config-tc-2)# l2-fwall
Zone1-SI-A(config-tc-2)# exit
```

The following commands add the firewalls. Three application ports (HTTP, FTP, and SNMP) are configured on each of the firewalls. The no-health-check parameter disables the Layer 4 health check for the specified application.

```
Zone1-SI-A(config)# server fw-name fw1 10.10.1.1
Zone1-SI-A(config-rs-fw1)# port http
Zone1-SI-A(config-rs-fw1)# port http no-health-check
Zone1-SI-A(config-rs-fw1)# port snmp
Zone1-SI-A(config-rs-fw1)# port snmp no-health-check
Zone1-SI-A(config-rs-fw1)# exit
Zone1-SI-A(config)# server fw-name fw2 10.10.1.2
Zone1-SI-A(config-rs-fw2)# port http
Zone1-SI-A(config-rs-fw2)# port http no-health-check
Zone1-SI-A(config-rs-fw2)# port snmp
Zone1-SI-A(config-rs-fw2)# port snmp no-health-check
Zone1-SI-A(config-rs-fw2)# exit
```

The following commands add the firewall definitions to the firewall port group (always group 2).

```
Zone1-SI-A(config)# server fw-group 2
Zone1-SI-A(config-tc-2)# fw-name fw1
Zone1-SI-A(config-tc-2)# fw-name fw2
```

The following command enables the active-active mode and specifies the priority of this ServerIron. In this case, ServerIron Zone1-SI-A has the higher priority. Its partner, ServerIron Zone1-SI-B, will be configured with a lower priority (1).

```
Zone1-SI-A(config-tc-2)# sym-priority 255
```

The following commands add the paths through the firewalls to the ServerIrons in zones 2 and 3. In addition, static MAC entries are added for the firewall interfaces.

NOTE: The path IDs must be in contiguous, ascending numerical order, starting with 1. For example, path sequence 1, 2, 3, 4 is valid. Path sequence 4, 3, 2, 1 or 1, 3, 4, 5 is not valid.

```
Zone1-SI-A(config-tc-2)# fwall-info 1 4/1 10.10.2.222 10.10.1.1
Zone1-SI-A(config-tc-2)# fwall-info 2 4/11 10.10.2.222 10.10.1.2
Zone1-SI-A(config-tc-2)# fwall-info 3 4/1 10.10.2.223 10.10.1.1
Zone1-SI-A(config-tc-2)# fwall-info 4 4/11 10.10.2.223 10.10.1.2
Zone1-SI-A(config-tc-2)# fwall-info 5 4/1 10.10.3.111 10.10.1.1
Zone1-SI-A(config-tc-2)# fwall-info 6 4/11 10.10.3.111 10.10.1.2
Zone1-SI-A(config-tc-2)# exit
```


The following commands set the load balancing method to balance requests based on the firewall that has the least number of connections for the requested service. For example, the ServerIron will load balance HTTP requests based on the firewall that has fewer HTTP session entries in the ServerIron session table.

```
Zone1-SI-A(config-tc-2)# fw-predictor per-service-least-conn
Zone1-SI-A(config-tc-2)# exit
```

The following commands configure standard IP ACLs for the IP sub-nets in one of the zones this ServerIron is not in.

```
Zone1-SI-A(config)# access-list 2 permit 10.10.2.0 0.0.0.255
Zone1-SI-A(config)# access-list 2 permit 10.10.8.0 0.0.0.255
```

The following commands configure the zone parameters. To configure a zone, specify a name for the zone, then a zone number (from 1 – 10), followed by the number of the ACL that specifies the IP addresses in the zone. In this example, the ACL numbers and zone numbers are the same, but this is not required.

```
Zone1-SI-A(config)# server fw-group 2
Zone1-SI-A(config-tc-2)# fwall-zone Zone2 2 2
Zone1-SI-A(config-tc-2)# exit
```

The following commands configure the SLB information. Each of the servers in zones 2 and 3 is added as a real server, then the servers are bound to a VIP. The servers are added using the **server remote-name** command instead of the **server real-name** command because the servers are not directly connected to the ServerIron. Instead, they are connected to the ServerIron through other routers (in this case, the firewalls).

```
Zone1-SI-A(config)# server remote-name web1 10.10.8.40
Zone1-SI-A(config-rs-web1)# port http
Zone1-SI-A(config-rs-web1)# exit
Zone1-SI-A(config)# server remote-name web2 10.10.8.42
Zone1-SI-A(config-rs-web2)# port http
Zone1-SI-A(config-rs-web2)# exit
Zone1-SI-A(config)# server remote-name web3 10.10.6.41
Zone1-SI-A(config-rs-web3)# port http
Zone1-SI-A(config-rs-web3)# exit
Zone1-SI-A(config)# server remote-name web4 10.10.6.43
Zone1-SI-A(config-rs-web4)# port http
Zone1-SI-A(config-rs-web4)# exit
Zone1-SI-A(config)# server virtual www.web.com 10.10.1.10
Zone1-SI-A(config-vs-www.web.com)# port http
Zone1-SI-A(config-vs-www.web.com)# bind http web1 http web2 http web3 http web4 http
Zone1-SI-A(config-vs-www.web.com)# exit
```

The following command enables SLB-to-FWLB.

```
Zone1-SI-A(config)# server slb-fw
```

The following commands enable FWLB.

```
Zone1-SI-A(config)# ip l4-policy 1 fw tcp 0 global
Zone1-SI-A(config)# ip l4-policy 2 fw udp 0 global
```

The following command saves the configuration changes to the startup-config file.

```
Zone1-SI-A(config)# write memory
```

Commands on Zone 1's Standby ServerIron (Zone1-SI-S)

```
ServerIron> enable
ServerIron# configure terminal
ServerIron(config)# hostname Zone1-SI-S
Zone1-SI-S(config)# vlan 1
Zone1-SI-S(config-vlan-1)# always-active
Zone1-SI-S(config-vlan-1)# exit
Zone1-SI-S(config)# vlan 2
Zone1-SI-S(config-vlan-2)# always-active
Zone1-SI-S(config-vlan-2)# tagged ethernet 4/11 to 4/12
```



```
Zone1-SI-S(config-vlan-2)# untagged ethernet 4/1 to 4/8
Zone1-SI-S(config-vlan-2)# router-interface ve 1
Zone1-SI-S(config-vlan-2)# exit
Zone1-SI-S(config)# interface ve 1
Zone1-SI-S(config-ve-1)# ip address 10.10.1.112 255.255.255.0
Zone1-SI-S(config-ve-1)# exit
Zone1-SI-S(config)# vlan 20
Zone1-SI-S(config-vlan-20)# always-active
Zone1-SI-S(config-vlan-20)# tagged ethernet 4/11 to 4/12
Zone1-SI-S(config-vlan-20)# untagged ethernet 4/13 to 4/24
Zone1-SI-S(config-vlan-20)# router-interface ve 2
Zone1-SI-S(config-vlan-20# exit
Zone1-SI-S(config)# interface ve 2
Zone1-SI-S(config-ve-2)# ip address 10.10.7.102 255.255.255.0
Zone1-SI-S(config-ve-2)# exit
Zone1-SI-S(config)# ip route 0.0.0.0 0.0.0.0 10.10.1.2
Zone1-SI-S(config)# ip route 20.20.0.0 255.255.0.0 10.10.7.100
Zone1-SI-S(config)# vlan 10
Zone1-SI-S(config-vlan-10)# untagged ethernet 4/9 to 4/10
Zone1-SI-S(config-vlan-10)# exit
Zone1-SI-S(config)# trunk switch ethernet 4/9 to 4/10
Zone1-SI-S(config)# server fw-port 4/9
Zone1-SI-S(config)# trunk switch ethernet 4/11 to 4/12
Zone1-SI-S(config)# server partner-ports ethernet 4/11
Zone1-SI-S(config)# server partner-ports ethernet 4/12
Zone1-SI-S(config)# server fw-group 2
Zone1-SI-S(config-tc-2)# l2-fwall
Zone1-SI-S(config-tc-2)# exit
Zone1-SI-S(config)# server fw-name fw1 10.10.1.1
Zone1-SI-S(config-rs-fw1)# port http
Zone1-SI-S(config-rs-fw1)# port http no-health-check
Zone1-SI-S(config-rs-fw1)# port snmp
Zone1-SI-S(config-rs-fw1)# port snmp no-health-check
Zone1-SI-S(config-rs-fw1)# exit
Zone1-SI-S(config)# server fw-name fw2 10.10.1.2
Zone1-SI-S(config-rs-fw2)# port http
Zone1-SI-S(config-rs-fw2)# port http no-health-check
Zone1-SI-S(config-rs-fw2)# port snmp
Zone1-SI-S(config-rs-fw2)# port snmp no-health-check
Zone1-SI-S(config-rs-fw2)# exit
Zone1-SI-S(config)# server fw-group 2
Zone1-SI-S(config-tc-2)# fw-name fw1
Zone1-SI-S(config-tc-2)# fw-name fw2
Zone1-SI-S(config-tc-2)# sym-priority 1
Zone1-SI-S(config-tc-2)# fwall-info 1 4/11 10.10.2.222 10.10.1.1
Zone1-SI-S(config-tc-2)# fwall-info 2 4/1 10.10.2.222 10.10.1.2
Zone1-SI-S(config-tc-2)# fwall-info 3 4/11 10.10.2.223 10.10.1.1
Zone1-SI-S(config-tc-2)# fwall-info 4 4/1 10.10.2.223 10.10.1.2
Zone1-SI-S(config-tc-2)# fwall-info 5 4/11 10.10.3.111 10.10.1.1
Zone1-SI-S(config-tc-2)# fwall-info 6 4/1 10.10.3.111 10.10.1.2
Zone1-SI-S(config-tc-2)# exit
Zone1-SI-S(config-tc-2)# fw-predictor per-service-least-conn
Zone1-SI-S(config-tc-2)# exit
Zone1-SI-S(config)# access-list 2 permit 10.10.2.0 0.0.0.255
Zone1-SI-S(config)# access-list 2 permit 10.10.8.0 0.0.0.255
Zone1-SI-S(config)# server fw-group 2
Zone1-SI-S(config-tc-2)# fwall-zone Zone2 2 2
Zone1-SI-S(config-tc-2)# exit
Zone1-SI-S(config)# server remote-name web1 10.10.8.40
```



```
Zone1-SI-S(config-rs-web1)# port http
Zone1-SI-S(config-rs-web1)# exit
Zone1-SI-S(config)# server remote-name web2 10.10.8.42
Zone1-SI-S(config-rs-web2)# port http
Zone1-SI-S(config-rs-web2)# exit
Zone1-SI-S(config)# server remote-name web3 10.10.6.41
Zone1-SI-S(config-rs-web3)# port http
Zone1-SI-S(config-rs-web3)# exit
Zone1-SI-S(config)# server remote-name web4 10.10.6.43
Zone1-SI-S(config-rs-web4)# port http
Zone1-SI-S(config-rs-web4)# exit
Zone1-SI-S(config)# server virtual www.web.com 10.10.1.10
Zone1-SI-S(config-vs-www.web.com)# port http
Zone1-SI-S(config-vs-www.web.com)# bind http web1 http web2 http web3 http web4 http
Zone1-SI-S(config-vs-www.web.com)# exit
Zone1-SI-S(config)# server slb-fw
Zone1-SI-S(config)# ip l4-policy 1 fw tcp 0 global
Zone1-SI-S(config)# ip l4-policy 2 fw udp 0 global
Zone1-SI-S(config)# write memory
```

Commands on Zone 2's Active ServerIron (Zone2-SI-A)

The following commands configure ServerIron Zone2-SI-A in zone 2. The configuration is similar to the configuration for ServerIron Zone1-SI-A, except the ACL and zone information are for zone 3, and FWLB-to-SLB is enabled instead of SLB-to-FWLB.

```
ServerIron> enable
ServerIron# configure terminal
ServerIron(config)# hostname Zone2-SI-A
Zone2-SI-A(config)# vlan 1
Zone2-SI-A(config-vlan-1)# always-active
Zone1-SI-A(config)# vlan 2
Zone1-SI-A(config-vlan-2)# always-active
Zone1-SI-A(config-vlan-2)# tagged ethernet 4/11 to 4/12
Zone1-SI-A(config-vlan-2)# untagged ethernet 4/1 to 4/8
Zone1-SI-A(config-vlan-2)# router-interface ve 1
Zone1-SI-A(config-vlan-2)# exit
Zone1-SI-A(config)# interface ve 1
Zone1-SI-A(config-ve-1)# ip address 10.10.2.222 255.255.255.0
Zone1-SI-A(config-ve-1)# exit
Zone1-SI-A(config)# vlan 20
Zone1-SI-A(config-vlan-20)# always-active
Zone1-SI-A(config-vlan-20)# tagged ethernet 4/11 to 4/12
Zone1-SI-A(config-vlan-20)# untagged ethernet 4/13 to 4/24
Zone1-SI-A(config-vlan-20)# router-interface ve 2
Zone1-SI-A(config-vlan-20# exit
Zone1-SI-A(config)# interface ve 2
Zone1-SI-A(config-ve-2)# ip address 10.10.8.101 255.255.255.0
Zone1-SI-A(config-ve-2)# exit
Zone2-SI-A(config)# ip route 0.0.0.0 0.0.0.0 10.10.2.1
Zone2-SI-A(config)# vlan 10
Zone2-SI-A(config-vlan-10)# untagged ethernet 4/9 to 4/10
Zone2-SI-A(config-vlan-10)# exit
Zone2-SI-A(config)# trunk switch ethernet 4/9 to 4/10
Zone2-SI-A(config)# server fw-port 4/9
Zone2-SI-A(config)# trunk switch ethernet 4/11 to 4/12
Zone2-SI-A(config)# server partner-ports ethernet 4/11
Zone2-SI-A(config)# server partner-ports ethernet 4/12
Zone2-SI-A(config)# server fw-group 2
```



```
Zone2-SI-A(config-tc-2)# l2-fwall
Zone2-SI-A(config-tc-2)# exit
Zone2-SI-A(config)# server fw-name fw1 10.10.2.1
Zone2-SI-A(config-rs-fw1)# port http
Zone2-SI-A(config-rs-fw1)# port http no-health-check
Zone2-SI-A(config-rs-fw1)# port ftp
Zone2-SI-A(config-rs-fw1)# port ftp no-health-check
Zone2-SI-A(config-rs-fw1)# port snmp
Zone2-SI-A(config-rs-fw1)# port snmp no-health-check
Zone2-SI-A(config-rs-fw1)# exit
Zone2-SI-A(config)# server fw-name fw2 10.10.2.2
Zone2-SI-A(config-rs-fw2)# port http
Zone2-SI-A(config-rs-fw2)# port http no-health-check
Zone2-SI-A(config-rs-fw2)# port ftp
Zone2-SI-A(config-rs-fw2)# port ftp no-health-check
Zone2-SI-A(config-rs-fw2)# port snmp
Zone2-SI-A(config-rs-fw2)# port snmp no-health-check
Zone2-SI-A(config-rs-fw2)# exit
Zone2-SI-A(config)# server fw-group 2
Zone2-SI-A(config-tc-2)# fw-name fw1
Zone2-SI-A(config-tc-2)# fw-name fw2
Zone2-SI-A(config-tc-2)# sym-priority 255
Zone2-SI-A(config-tc-2)# fwall-info 1 4/1 10.10.1.111 10.10.2.1
Zone2-SI-A(config-tc-2)# fwall-info 2 4/11 10.10.1.111 10.10.2.2
Zone2-SI-A(config-tc-2)# fwall-info 3 4/1 10.10.1.112 10.10.2.1
Zone2-SI-A(config-tc-2)# fwall-info 4 4/11 10.10.1.112 10.10.2.2
Zone2-SI-A(config-tc-2)# fwall-info 5 4/1 10.10.3.111 10.10.2.1
Zone2-SI-A(config-tc-2)# fwall-info 6 4/11 10.10.3.111 10.10.2.2
Zone2-SI-A(config-tc-2)# exit
Zone2-SI-A(config)# server fw-group 2
Zone2-SI-A(config-tc-2)# fw-predictor per-service-least-conn
Zone2-SI-A(config-tc-2)# exit
Zone2-SI-A(config)# access-list 3 permit 10.10.3.0 0.0.0.255
Zone2-SI-A(config)# access-list 3 permit 10.10.6.0 0.0.0.255
Zone2-SI-A(config)# server fw-group 2
Zone2-SI-A(config-tc-2)# fwall-zone zone3 3 3
Zone2-SI-A(config-tc-2)# exit
Zone2-SI-A(config)# server real-name rs1 10.10.8.40
Zone2-SI-A(config-rs-rs1)# port http
Zone2-SI-A(config-rs-rs1)# exit
Zone2-SI-A(config)# server real-name rs1 10.10.8.42
Zone2-SI-A(config-rs-rs2)# port http
Zone2-SI-A(config-rs-rs2)# exit
Zone2-SI-A(config)# server virtual www.rs.com 10.10.8.10
Zone2-SI-A(config-vs-www.rs.com)# port http
Zone2-SI-A(config-vs-www.web.com)# bind http rs1 http rs2 http
Zone2-SI-A(config-vs-www.web.com)# exit
Zone2-SI-A(config)# server fw-slb
Zone2-SI-A(config)# ip l4-policy 1 fw tcp 0 global
Zone2-SI-A(config)# ip l4-policy 2 fw udp 0 global
Zone2-SI-A(config)# write memory
```

Commands on Zone 2's Standby ServerIron (Zone2-SI-S)

```
ServerIron> enable
ServerIron# configure terminal
ServerIron(config)# hostname Zone2-SI-S
Zone2-SI-S(config)# vlan 1
Zone2-SI-S(config-vlan-1)# always-active
```



```
Zone1-SI-S(config)# vlan 2
Zone1-SI-S(config-vlan-2)# always-active
Zone1-SI-S(config-vlan-2)# tagged ethernet 4/11 to 4/12
Zone1-SI-S(config-vlan-2)# untagged ethernet 4/1 to 4/8
Zone1-SI-S(config-vlan-2)# router-interface ve 1
Zone1-SI-S(config-vlan-2)# exit
Zone1-SI-S(config)# interface ve 1
Zone1-SI-S(config-ve-1)# ip address 10.10.2.223 255.255.255.0
Zone1-SI-S(config-ve-1)# exit
Zone1-SI-S(config)# vlan 20
Zone1-SI-S(config-vlan-20)# always-active
Zone1-SI-S(config-vlan-20)# tagged ethernet 4/11 to 4/12
Zone1-SI-S(config-vlan-20)# untagged ethernet 4/13 to 4/24
Zone1-SI-S(config-vlan-20)# router-interface ve 2
Zone1-SI-S(config-vlan-20# exit
Zone1-SI-S(config)# interface ve 2
Zone1-SI-S(config-ve-2)# ip address 10.10.8.102 255.255.255.0
Zone1-SI-S(config-ve-2)# exit
Zone2-SI-S(config)# ip route 0.0.0.0 0.0.0.0 10.10.2.2
Zone2-SI-S(config)# vlan 10
Zone2-SI-S(config-vlan-10)# untagged ethernet 4/9 to 4/10
Zone2-SI-S(config-vlan-10)# exit
Zone2-SI-S(config)# trunk switch ethernet 4/9 to 4/10
Zone2-SI-S(config)# server fw-port 4/9
Zone2-SI-S(config)# trunk switch ethernet 4/11 to 4/12
Zone2-SI-S(config)# server partner-ports ethernet 4/11
Zone2-SI-S(config)# server partner-ports ethernet 4/12
Zone2-SI-S(config)# server fw-group 2
Zone2-SI-S(config-tc-2)# l2-fwall
Zone2-SI-S(config-tc-2)# exit
Zone2-SI-S(config)# server fw-name fw1 10.10.2.1
Zone2-SI-S(config-rs-fw1)# port http
Zone2-SI-S(config-rs-fw1)# port http no-health-check
Zone2-SI-S(config-rs-fw1)# port ftp
Zone2-SI-S(config-rs-fw1)# port ftp no-health-check
Zone2-SI-S(config-rs-fw1)# port snmp
Zone2-SI-S(config-rs-fw1)# port snmp no-health-check
Zone2-SI-S(config-rs-fw1)# exit
Zone2-SI-S(config)# server fw-name fw2 10.10.2.2
Zone2-SI-S(config-rs-fw2)# port http
Zone2-SI-S(config-rs-fw2)# port http no-health-check
Zone2-SI-S(config-rs-fw2)# port ftp
Zone2-SI-S(config-rs-fw2)# port ftp no-health-check
Zone2-SI-S(config-rs-fw2)# port snmp
Zone2-SI-S(config-rs-fw2)# port snmp no-health-check
Zone2-SI-S(config-rs-fw2)# exit
Zone2-SI-S(config)# server fw-group 2
Zone2-SI-S(config-tc-2)# fw-name fw1
Zone2-SI-S(config-tc-2)# fw-name fw2
Zone2-SI-S(config-tc-2)# sym-priority 1
Zone2-SI-S(config-tc-2)# fwall-info 1 4/11 10.10.1.111 10.10.2.1
Zone2-SI-S(config-tc-2)# fwall-info 2 4/1 10.10.1.111 10.10.2.2
Zone2-SI-S(config-tc-2)# fwall-info 3 4/11 10.10.1.112 10.10.2.1
Zone2-SI-S(config-tc-2)# fwall-info 4 4/1 10.10.1.112 10.10.2.2
Zone2-SI-S(config-tc-2)# fwall-info 5 4/11 10.10.3.111 10.10.2.1
Zone2-SI-S(config-tc-2)# fwall-info 6 4/1 10.10.3.111 10.10.2.2
Zone2-SI-S(config-tc-2)# exit
Zone2-SI-S(config)# server fw-group 2
Zone2-SI-S(config-tc-2)# fw-predictor per-service-least-conn
```



```
Zone2-SI-S(config-tc-2)# exit
Zone2-SI-S(config)# access-list 3 permit 10.10.3.0 0.0.0.255
Zone2-SI-S(config)# access-list 3 permit 10.10.6.0 0.0.0.255
Zone2-SI-S(config)# server fw-group 2
Zone2-SI-S(config-tc-2)# fwall-zone zone3 3 3
Zone2-SI-S(config-tc-2)# exit
Zone2-SI-S(config)# server real-name rs1 10.10.8.40
Zone2-SI-S(config-rs-rs1)# port http
Zone2-SI-S(config-rs-rs1)# exit
Zone2-SI-S(config)# server real-name rs1 10.10.8.42
Zone2-SI-S(config-rs-rs2)# port http
Zone2-SI-S(config-rs-rs2)# exit
Zone2-SI-S(config)# server virtual www.rs.com 10.10.8.10
Zone2-SI-S(config-vs-www.rs.com)# port http
Zone2-SI-S(config-vs-www.web.com)# bind http rs1 http rs2 http
Zone2-SI-S(config-vs-www.web.com)# exit
Zone2-SI-S(config)# server fw-slb
Zone2-SI-S(config)# ip l4-policy 1 fw tcp 0 global
Zone2-SI-S(config)# ip l4-policy 2 fw udp 0 global
Zone2-SI-S(config)# write memory
```

Commands on Zone 3's ServerIron (Zone3-SI-A)

Here are the commands for configuring the ServerIron in zone 3.

```
ServerIron> enable
ServerIron# configure terminal
ServerIron(config)# hostname Zone3-SI-A
Zone3-SI-A(config)# vlan 1
Zone3-SI-A(config-vlan-1)# untagged ethernet 4/1 to 4/12
Zone3-SI-A(config-vlan-1)# router-interface ve 1
Zone3-SI-A(config-vlan-1)# exit
Zone3-SI-A(config)# interface ve 1
Zone3-SI-A(config-ve-1)# ip address 10.10.3.111 255.255.255.0
Zone3-SI-A(config-ve-1)# exit
Zone3-SI-A(config)# vlan 2
Zone3-SI-A(config-vlan-2)# untagged ethernet 4/13 to 4/24
Zone3-SI-A(config-vlan-2)# router-interface ve 2
Zone3-SI-A(config-vlan-2)# exit
Zone3-SI-A(config)# interface ve 2
Zone3-SI-A(config-ve-1)# ip address 10.10.6.101 255.255.255.0
Zone3-SI-A(config-ve-1)# exit
Zone3-SI-A(config)# ip route 0.0.0.0 0.0.0.0 10.10.3.1
Zone3-SI-A(config)# server fw-name fw1 10.10.3.1
Zone3-SI-A(config-rs-fw1)# port http
Zone3-SI-A(config-rs-fw1)# port http no-health-check
Zone3-SI-A(config-rs-fw1)# port ftp
Zone3-SI-A(config-rs-fw1)# port ftp no-health-check
Zone3-SI-A(config-rs-fw1)# port snmp
Zone3-SI-A(config-rs-fw1)# port snmp no-health-check
Zone3-SI-A(config-rs-fw1)# exit
Zone3-SI-A(config)# server fw-name fw2 10.10.3.2
Zone3-SI-A(config-rs-fw2)# port http
Zone3-SI-A(config-rs-fw2)# port http no-health-check
Zone3-SI-A(config-rs-fw2)# port ftp
Zone3-SI-A(config-rs-fw2)# port ftp no-health-check
Zone3-SI-A(config-rs-fw2)# port snmp
Zone3-SI-A(config-rs-fw2)# port snmp no-health-check
Zone3-SI-A(config-rs-fw2)# exit
```



```
Zone3-SI-A(config)# server fw-group 2
Zone3-SI-A(config-tc-2)# fw-name fw1
Zone3-SI-A(config-tc-2)# fw-name fw2
Zone3-SI-A(config-tc-2)# fwall-info 1 4/1 10.10.1.111 10.10.3.1
Zone3-SI-A(config-tc-2)# fwall-info 2 4/2 10.10.1.111 10.10.3.2
Zone3-SI-A(config-tc-2)# fwall-info 3 4/1 10.10.1.112 10.10.3.1
Zone3-SI-A(config-tc-2)# fwall-info 4 4/2 10.10.1.112 10.10.3.2
Zone3-SI-A(config-tc-2)# fwall-info 5 4/1 10.10.2.222 10.10.3.1
Zone3-SI-A(config-tc-2)# fwall-info 6 4/2 10.10.2.222 10.10.3.2
Zone3-SI-A(config-tc-2)# fwall-info 7 4/1 10.10.2.223 10.10.3.1
Zone3-SI-A(config-tc-2)# fwall-info 8 4/2 10.10.2.223 10.10.3.2
Zone3-SI-A(config-tc-2)# exit
Zone3-SI-A(config)# server fw-group 2
Zone3-SI-A(config-tc-2)# fw-predictor per-service-least-conn
Zone3-SI-A(config-tc-2)# exit
Zone3-SI-A(config)# access-list 2 permit 10.10.2.0 0.0.0.255
Zone3-SI-A(config)# access-list 2 permit 10.10.8.0 0.0.0.255
Zone3-SI-A(config)# server fw-group 2
Zone3-SI-A(config-tc-2)# fwall-zone zone2 2 2
Zone3-SI-A(config-tc-2)# exit
Zone3-SI-A(config)# server real-name sr1 10.10.6.41
Zone3-SI-A(config-rs-sr1)# port http
Zone3-SI-A(config-rs-sr1)# exit
Zone3-SI-A(config)# server real-name sr2 10.10.6.43
Zone3-SI-A(config-rs-sr2)# port http
Zone3-SI-A(config-rs-sr2)# exit
Zone3-SI-A(config)# server virtual www.sr.com 10.10.6.10
Zone3-SI-A(config-vs-www.rs.com)# port http
Zone3-SI-A(config-vs-www.web.com)# bind http sr2 http sr1 http
Zone3-SI-A(config-vs-www.web.com)# exit
Zone3-SI-A(config)# server fw-slb
Zone3-SI-A(config)# ip l4-policy 1 fw tcp 0 global
Zone3-SI-A(config)# ip l4-policy 2 fw udp 0 global
Zone3-SI-A(config)# write memory
```

Chapter 7
Configuring FWLB for NAT Firewalls

Some Layer 3 firewalls perform network address translation (NAT). These firewalls translate private addresses on the private side of the network into public (Internet) addresses on the public side of the network.

NOTE: The configuration steps for firewalls that perform NAT are identical to the steps for basic and IronClad FWLB without NAT, with just one additional step. The additional step disables load balancing for the NAT addresses. The following sections provide more information.

You can deploy ServerIrons to load balance NAT firewalls in a basic configuration or an IronClad configuration, just as in the examples in the previous sections. Configuring the ServerIrons for NAT requires only one additional step. The additional step disables load balancing for the NAT addresses, which are the addresses the firewalls use when translating private addresses into Internet addresses.

You can configure a single ServerIron on each side of the firewalls (as in the basic configuration example in Figure 7.1) or you can configure active-standby pairs of ServerIrons on each side of the firewalls (as in Figure 7.2).

Firewalls perform NAT in either of the following ways. The ServerIron supports load balancing for either method and the ServerIron configuration is the same for each method. You do not need to know which method your firewalls are using to configure the ServerIrons to load balance for them.

- Hiding internal addresses behind a single public address – The firewall is configured with a single Internet address that it uses for clients that initiate traffic from within the private side of the network. The firewall translates the source address for such traffic from the private address of the client into the public address. The firewall keeps track of the private addresses by including a Layer 4 port number from a pool of such numbers. When the firewall receives a return packet from a destination, the firewall uses the port number to identify the correct private address and translates the packet's destination address from the public address into the correct private address.
- Static translation – For traffic from a client inside the private network to a destination on the Internet, the firewall translates the private address into a unique Internet address. Likewise, for traffic from the Internet, the firewall translates the public address into a private address. Unlike the method above, the static method assigns a different, unique Internet address for each client in the private network. The method above uses a common Internet address for all private addresses.

Configuring Basic Layer 3 FWLB for NAT Firewalls

Figure 7.1 shows an example of a basic FWLB configuration for Layer 3 NAT firewalls. The procedures and CLI configuration example in this section are based on this sample configuration.

NOTE: The configuration steps for firewalls that perform NAT are identical to the steps for basic and IronClad FWLB without NAT, with just one additional step. The additional step disables load balancing for the NAT addresses. See “Preventing Load Balancing of the NAT Addresses” on page 7-5.

Figure 7.1 FWLB for Layer 3 firewalls performing NAT—basic configuration

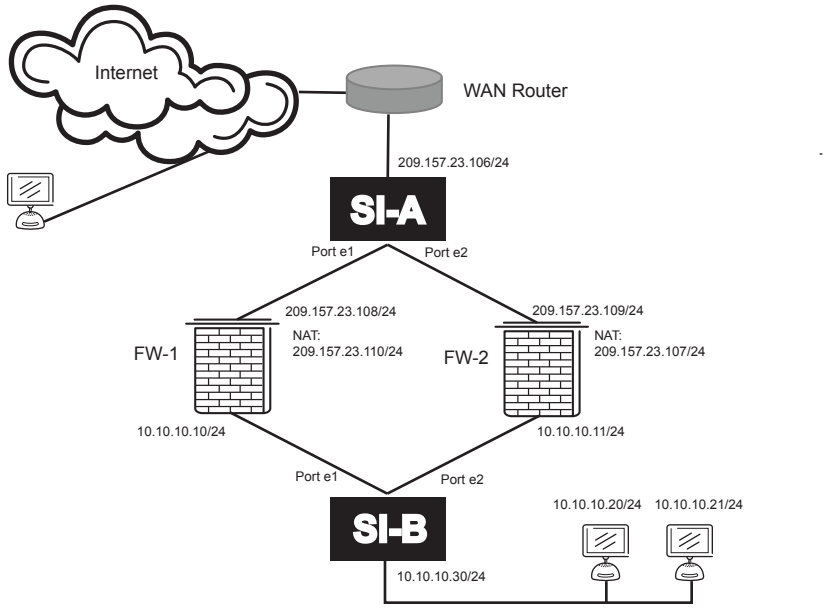

To configure basic Layer 3 FWLB for NAT firewalls, perform the following tasks.

Table 7.1: Configuration tasks – Basic FWLB for NAT Firewalls

Task	See page...
Configure Global Parameters	
Globally enable FWLB	7-2
Configure Firewall Parameters	
Define the firewalls and add them to the firewall group	7-3
Configure Firewall Group Parameters	
Configure the paths and add static MAC entries for the firewall interfaces with the ServerIron	7-4
Configure NAT Address Parameters	
Disable load balancing for the NAT addresses	7-5

Enabling FWLB

To enable FWLB, you configure global IP policies. FWLB for TCP and UDP is controlled independently, so you need to configure a separate global IP policy for each type of traffic.

When you enable FWLB for TCP or UDP globally, all ports that are in the firewall group are enabled for FWLB. All ServerIron ports are in firewall group 2 by default. Thus, if you enable FWLB globally, it affects all physical ports unless you remove ports from firewall groups.

NOTE: The user interface allows you to enable FWLB locally instead of globally. However, local policies are not applicable to FWLB. Enable the feature globally.

To enable FWLB globally, use the following method.

USING THE CLI

Enter the following commands at the global CONFIG level to enable FWLB for all TCP and UDP traffic:

```
ServerIron(config)# ip policy 1 fw tcp 0 global
ServerIron(config)# ip policy 2 fw udp 0 global
```

Syntax: [no] ip policy <policy-num> fw tcp | udp 0 global

The <policy-num> value identifies the policy and can be a number from 1 – 64.

Each policy affects TCP or UDP traffic, so you must specify **tcp** or **udp**.

The value 0 following the **tcp | udp** parameter specifies that the policy applies to all ports of the specified type (TCP or UDP). In this command, "0" is equivalent to "any port number". For FWLB, you must specify "0".

NOTE: Generally, the firewall itself performs validation and authentication for the traffic, so allowing the ServerIron to pass all traffic of the specified type (TCP or UDP) to the firewall simplifies configuration.

Defining the Firewalls and Adding Them to the Firewall Group

When FWLB is enabled, all the ServerIron ports are in firewall group 2 by default. However, you need to add an entry for each firewall, then add the firewalls to the firewall group. To add an entry for a firewall, specify the firewall name and IP address. You can specify a name up to 32 characters long.

NOTE: When static or NAT is used on firewalls in FWLB configurations, ServerIrons' virtual routing interface IP addresses that are in firewalls subnets should be excluded from NAT translation; otherwise there will be problems with firewall paths

To define the firewalls shown in Figure 7.1, use the following method.

USING THE CLI

To define the firewalls using the CLI, enter the following commands:

Commands for ServerIron A (External)

```
ServerIron-A(config)# server fw-name fw1 209.157.23.108
ServerIron-A(config-rs-fw1)# exit
ServerIron-A(config)# server fw-name fw2 209.157.23.109
ServerIron-A(config-rs-fw2)# exit
ServerIron-A(config)# server fw-group 2
ServerIron-A(config-tc-2)# fw-name fw1
ServerIron-A(config-tc-2)# fw-name fw254
```

Commands for ServerIron B (Internal)

```
ServerIron-B(config)# server fw-name fw1 10.10.10.10
ServerIron-B(config-rs-fw1)# exit
ServerIron-B(config)# server fw-name fw2 10.10.10.11
ServerIron-B(config-rs-fw2)# exit
ServerIron-B(config)# server fw-group 2
ServerIron-B(config-tc-2)# fw-name fw1
ServerIron-B(config-tc-2)# fw-name fw2
```

Command Syntax

Syntax: [no] server fw-name <string> <ip-addr>

NOTE: When you add a firewall name, the CLI level changes to the Firewall level. This level is used when you are configuring stateful FWLB.

Syntax: server fw-group 2

This command changes the CLI to firewall group configuration level. The firewall group number is 2. Only one firewall group is supported.

Syntax: [no] fw-name <string>

Adds a configured firewall to the firewall group.

Configuring the Paths and Adding Static MAC Entries

A path is configuration information the ServerIron uses to ensure that a given source and destination IP pair is always authenticated by the same Layer 3 firewall.

Each path consists of the following parameters:

- The path ID – A number that identifies the path. The paths go from one ServerIron to the other through the firewalls.
- The ServerIron port – The number of the port that connects the ServerIron to the firewall.
- The other ServerIron's or Layer 2 switch's IP address – The management address of the ServerIron or Layer 2 switch on the other side of the firewall. The ServerIron on the private network side and the other ServerIron or Layer 2 switch are the end points of the data path through the firewall.
- The next-hop IP address – The IP address of the firewall interface connected to this ServerIron.

For each type of firewall (Layer 3 synchronous and asynchronous, with or without NAT), you must configure paths between the ServerIrons through the firewalls.

In addition to configuring the paths, you need to create a static MAC entry for each firewall interface attached to the ServerIron.

NOTE: FWLB paths must be fully meshed. When you configure a FWLB path on a ServerIron, make sure you also configure a reciprocal path on the ServerIron attached to the other end of the firewalls. For example, if you configure four paths to four separate firewalls, make sure you configure four paths on the other ServerIron.

NOTE: The static MAC entries are required. You must add a static MAC entry for each firewall interface with the ServerIron.

To configure a path and add static MAC entries, use one of the following methods.

USING THE CLI

To configure the paths and static MAC entries for the configuration shown in Figure 3.2 on page 3-7, enter the following commands. Enter the first group of commands on ServerIron A. Enter the second group of commands on ServerIron B.

Commands for ServerIron A (External)

```
ServerIron-A(config)# server fw-group 2
ServerIron-A(config-tc-2)# fwall-info 1 1 10.10.10.30 209.157.23.108
ServerIron-A(config-tc-2)# fwall-info 2 2 10.10.10.30 209.157.23.109
ServerIron-A(config-tc-2)# exit
ServerIron-A(config)# static-mac-address abcd.da10.dc2c ethernet 1 high-priority
router-type
ServerIron-A(config)# static-mac-address abcd.da10.dc3f ethernet 2 high-priority
```



```
router-type
```

Commands for ServerIron B (Internal)

```
ServerIron-B(config)# server fw-group 2
ServerIron-B(config-tc-2)# fwall-info 1 1 209.157.23.106 10.10.10.10
ServerIron-B(config-tc-2)# fwall-info 2 2 209.157.23.106 10.10.10.11
ServerIron-B(config-tc-2)# exit
ServerIron-B(config)# static-mac-address abcd.da68.6655 ethernet 1 high-priority
router-type
ServerIron-B(config)# static-mac-address abcd.da68.6104 ethernet 2 high-priority
router-type
```

Command Syntax

Syntax: server fw-group 2

Syntax: [no] fwall-info <path-num> <portnum> <other-ServerIron-ip> <next-hop-ip>

The syntax for adding static MAC entries differs depending on whether you are using a stackable or chassis ServerIron.

Syntax for chassis devices:

Syntax: [no] static-mac-address <mac-addr> ethernet <portnum> [priority <0-7>] [host-type | router-type]

Syntax for stackable devices:

Syntax: static-mac-address <mac-addr> ethernet <portnum> [to <portnum> ethernet <portnum>] [normal-priority | high-priority] [host-type | router-type | fixed-host]

The priority can be 0 – 7 (0 is lowest and 7 is highest) for chassis devices and either normal-priority or high-priority for stackable devices.

The defaults are **host-type** and **0** or **normal-priority**.

NOTE: The static MAC entries are required. You must add a static MAC entry for each firewall interface with the ServerIron. In addition, you must use the **high-priority** and **router-type** parameters with the **static-mac-address** command. These parameters enable the ServerIron to use the address for FWLB.

NOTE: If you enter the command at the global CONFIG level, the static MAC entry applies to the default port-based VLAN (VLAN 1). If you enter the command at the configuration level for a specific port-based VLAN, the entry applies to that VLAN and not to the default VLAN.

Preventing Load Balancing of the NAT Addresses

When you configure ServerIrons for load balancing traffic across NAT firewalls, you must disable load balancing on the NAT addresses themselves. You can use either of the following methods to do so. Each method is equally valid and only one of the methods is required. You need to use one of these methods only on the ServerIron connected to the external network, not the ServerIron on the internal side of the network.

- Configure the NAT addresses as firewall addresses, but do not configure paths for the addresses. (This is shown below in the "Extra Firewall Method" section.)
- Configure IP access policies (filters) to deny load balancing for traffic addressed to the NAT addresses. (This is shown below in the "Access Policy Method" section.)

NOTE: In FWLB configurations, the IP policies do not block traffic altogether. They deny load balancing for the traffic. Thus, the ServerIron does not load balance packets addressed to the NAT addresses, but instead sends the traffic only to the firewall that originally sent the traffic.

USING THE CLI

Use either of the following methods to disable load balancing for the NAT addresses.

Extra Firewall Method

To disable load balancing for the NAT addresses by adding firewalls for the addresses, enter commands such as the following.

NOTE: Do not configure paths for the firewalls.

```
ServerIron-A(config)# server fw-name fw3NAT 209.157.23.107
ServerIron-A(config-rs-fw3NAT)# exit
ServerIron-A(config)# server fw-name fw4NAT 209.157.23.110
ServerIron-A(config-rs-fw4NAT)# exit
```

Access Policy Method

To disable load balancing for the NAT addresses using IP access policies, enter commands such as the following.

```
ServerIron-A(config)# ip filter 1 deny any 209.157.23.110 255.255.255.255
ServerIron-A(config)# ip filter 2 deny any 209.157.23.107 255.255.255.255
ServerIron-A(config)# ip filter 1024 permit any any
```

The first two commands configure policies to deny load balancing for the two NAT addresses. The third command allows all other traffic to be load balanced.

NOTE: The third policy, which permits all traffic, is required because once you define an access policy, the default action for packets that do not match a policy is to deny them. Thus, if you configure only the first two policies and not the third one, you actually disable load balancing altogether by denying the load balancing for all packets.

Configuration Example for FWLB with Layer 3 NAT Firewalls

This section shows the CLI commands for implementing the configuration shown in Figure 7.1. Note that the configuration steps are similar to those required for the basic configuration shown in Figure 3.2 on page 3-7. The only additional step required is to ensure that the ServerIron connected to the external network does not load balance return traffic to the addresses the firewalls use for NAT. For example, ServerIron A in Figure 7.1 must be configured so that it does not load balance return traffic to 209.157.23.107/24 or 209.157.23.110/24.

CLI Commands on ServerIron A (External)

The following commands configure ServerIron-A in Figure 7.1 for FWLB.

The **hostname** command changes the host name of the device to match the name used in Figure 7.1. The **ip address** and **ip default-gateway** commands configure the device's management IP address and its default gateway.

The **no span** command disables the Spanning Tree Protocol (STP) on the ServerIron.

```
ServerIron(config)# hostname ServerIron-A
ServerIron-A(config)# ip address 209.157.23.106 255.255.255.0
ServerIron-A(config)# ip default-gateway 209.157.23.108
ServerIron-A(config)# no span
```

The following two commands add the firewalls. The IP addresses are the firewalls' interfaces with the ServerIron.


```
ServerIron-A(config)# server fw-name fw1 209.157.23.108
ServerIron-A(config-rs-fw1)# exit
ServerIron-A(config)# server fw-name fw2 209.157.23.109
ServerIron-A(config-rs-fw2)# exit
```

The following two commands add firewall entries for the hidden NAT addresses. These entries prevent the ServerIron from load balancing the firewall traffic to these addresses. The ServerIron forwards a return packet addressed to one of these firewalls directly to the firewall that sent it, instead of using the hash mechanism to select a path for the traffic.

```
ServerIron-A(config)# server fw-name fw3NAT 209.157.23.107
ServerIron-A(config-rs-fw3NAT)# exit
ServerIron-A(config)# server fw-name fw4NAT 209.157.23.110
ServerIron-A(config-rs-fw4NAT)# exit
```

The following commands configure the firewall group parameters. The first commands change the CLI to the firewall group configuration level. The **fw-name** commands add the firewalls. Notice that the firewall definitions created above for the two NAT addresses are not added.

The **fwall-info** commands add paths from this ServerIron to the other one through the firewalls. Notice that no paths are configured for the firewall definitions created for the NAT addresses.

The **fw-name** <firewall-name> command adds the firewalls to the firewall group.

```
ServerIron-A(config)# server fw-group 2
ServerIron-A(config-tc-2)# fw-name fw1
ServerIron-A(config-tc-2)# fw-name fw2
ServerIron-A(config-tc-2)# fwall-info 1 1 10.10.10.30 209.157.23.108
ServerIron-A(config-tc-2)# fwall-info 2 2 10.10.10.30 209.157.23.109
ServerIron-A(config-tc-2)# exit
```

The following commands enable FWLB. You must enter the commands exactly as shown.

```
ServerIron-A(config)# ip policy 1 fw tcp 0 global
ServerIron-A(config)# ip policy 2 fw udp 0 global
```

The following commands add static MAC entries for the firewalls' interfaces with the ServerIron. The **high-priority** and **router-type** parameters are required for FWLB with Layer 3 firewalls.

```
ServerIron-A(config)# static-mac-address abcd.da10.dc2c ethernet 1 high-priority
router-type
ServerIron-A(config)# static-mac-address abcd.da10.dc3f ethernet 2 high-priority
router-type
```

The **write memory** command saves the configuration changes to the ServerIron's startup-config file on the device's flash memory.

```
ServerIron-A(config)# write memory
```

Alternative Configuration for ServerIron A

The example above configures FWLB for NAT firewalls by adding firewall definitions for the IP addresses the NAT service on the firewalls uses for traffic sent from a client inside the firewalls to a destination outside the firewalls.

Alternatively, you can configure IP access policies that deny load balancing for the NAT addresses. For the example in Figure 7.1 on page 7-2, you would enter the following commands:

```
ServerIron-A(config)# ip filter 1 deny any 209.157.23.110 255.255.255.255
ServerIron-A(config)# ip filter 2 deny any 209.157.23.107 255.255.255.255
ServerIron-A(config)# ip filter 1024 permit any any
```

The first two commands configure policies to deny load balancing for the two NAT addresses. The third command allows all other traffic to be load balanced.

NOTE: The third policy, which permits all traffic, is required because once you define an access policy, the default action for packets that do not match a policy is to deny them. Thus, if you configure only the first two policies and not the third one, you actually disable load balancing altogether by denying the load balancing for all packets.

The other commands are the same as in the previous section.

CLI Commands on ServerIron B (Internal)

To following CLI commands configure ServerIron B in Figure 7.1. Notice that this ServerIron is not configured to deny load balancing for the NAT addresses used by the firewalls. This ServerIron sees only the internal addresses, not the NAT addresses.

```
ServerIron-B(config)# hostname ServerIron-B
ServerIron-B(config)# ip address 10.10.10.30 255.255.255.0
ServerIron-B(config)# ip default-gateway 10.10.10.10
ServerIron-B(config)# no span
ServerIron-B(config)# server fw-name fw1 10.10.10.10
ServerIron-B(config-rs-fw1)# exit
ServerIron-B(config)# server fw-name fw2 10.10.10.11
ServerIron-B(config-rs-fw2)# exit
ServerIron-B(config)# server fw-group 2
ServerIron-B(config-tc-2)# fw-name fw1
ServerIron-B(config-tc-2)# fw-name fw2
ServerIron-B(config-tc-2)# fwall-info 1 1 209.157.23.106 10.10.10.10
ServerIron-B(config-tc-2)# fwall-info 2 2 209.157.23.106 10.10.10.11
ServerIron-B(config-tc-2)# exit
ServerIron-B(config)# static-mac-address abcd.da68.6655 ethernet 1 high-priority
router-type
ServerIron-B(config)# static-mac-address abcd.da68.6104 ethernet 2 high-priority
router-type
ServerIron-B(config)# ip policy 1 fw tcp 0 global
ServerIron-B(config)# ip policy 2 fw udp 0 global
```

Configuring IronClad Layer 3 FWLB for NAT

Figure 7.2 shows an example of an IronClad FWLB configuration for Layer 3 NAT firewalls. The procedures and CLI configuration example in this section are based on this sample configuration.

NOTE: The configuration steps for firewalls that perform NAT are identical to the steps for basic and IronClad FWLB without NAT, with just one additional step. The additional step disables load balancing for the NAT addresses. See “Preventing Load Balancing of the NAT Addresses” on page 7-15.

Figure 7.2 FWLB for Layer 3 firewalls performing NAT—IronClad configuration

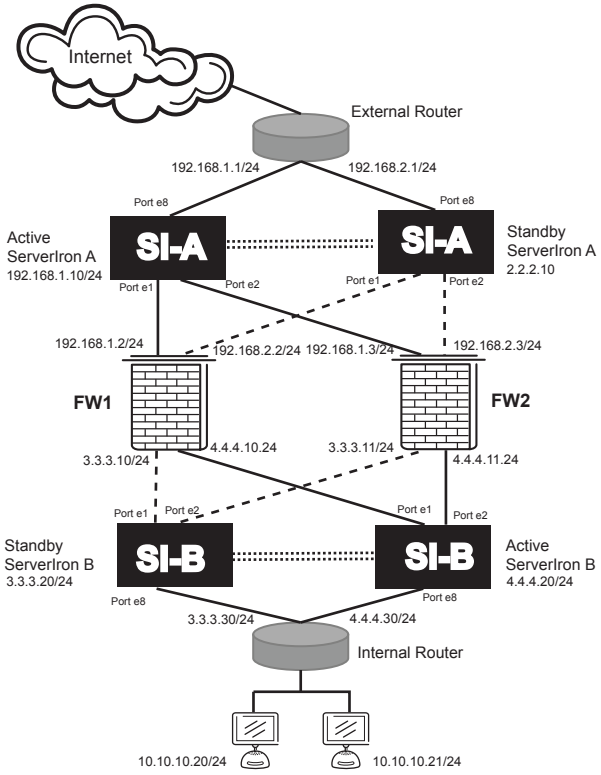

To configure IronClad FWLB for NAT firewalls, perform the following tasks.

Table 7.2: Configuration tasks – IronClad FWLB for NAT Firewalls

Task	See page...
Configure Global Parameters	
Globally enable FWLB	7-10
Identify the partner port (the link between the active and standby ServerIrons)	7-10
Identify the router port (ServerIron ports connected to routers)	7-10
Configure Firewall Parameters	
Define the firewalls and add them to the firewall group	7-11
Configure Firewall Group Parameters	
Configure the paths and add static MAC entries for the firewall interfaces with the ServerIron	7-12
Specify the ServerIron priority (determines which ServerIron in the active-standby pair is the default active ServerIron)	7-14
Configure NAT Address Parameters	
Disable load balancing for the NAT addresses	7-15

Enabling FWLB

To enable FWLB, you configure global IP policies. FWLB for TCP and UDP is controlled independently, so you need to configure a separate global IP policy for each type of traffic.

When you enable FWLB for TCP or UDP globally, all ports that are in the firewall group are enabled for FWLB. All ServerIron ports are in firewall group 2 by default. Thus, if you enable FWLB globally, it affects all physical ports unless you remove ports from firewall groups.

NOTE: The user interface allows you to enable FWLB locally instead of globally. However, local policies are not applicable to FWLB. Enable the feature globally.

To enable FWLB globally, use the following method.

USING THE CLI

Enter the following commands at the global CONFIG level to enable FWLB for all TCP and UDP traffic:

```
ServerIron(config)# ip policy 1 fw tcp 0 global
ServerIron(config)# ip policy 2 fw udp 0 global
```

Syntax: [no] ip policy <policy-num> fw tcp | udp 0 global

The <policy-num> value identifies the policy and can be a number from 1 – 64.

Each policy affects TCP or UDP traffic, so you must specify **tcp** or **udp**.

The value 0 following the **tcp | udp** parameter specifies that the policy applies to all ports of the specified type (TCP or UDP). In this command, “0” is equivalent to “any port number”. For FWLB, you must specify “0”.

NOTE: Generally, the firewall itself performs validation and authentication for the traffic, so allowing the ServerIron to pass all traffic of the specified type (TCP or UDP) to the firewall simplifies configuration.

Specifying the Partner Port

If you are configuring the ServerIron for IronClad FWLB, you need to specify the port number of the dedicated link between the ServerIron and its partner.

USING THE CLI

To specify the port, enter a command such as the following at the global CLI level:

```
ServerIron(config)# server fw-port 5
```

Syntax: [no] server fw-port <portnum>

If the link between the two ServerIrons is a trunk group (recommended for added redundancy), specify the port number of the primary port. The primary port is the first port in the trunk group.

Specifying the Router Ports

IronClad FWLB configurations require paths to the routers as part of the active-standby configuration for the ServerIrons. You need to identify the ports on the ServerIron that are attached to the router(s).

USING THE CLI

To identify port 8 on a ServerIron as a router port, enter the following command:

```
ServerIron(config)# server router-port 8
```

Syntax: [no] server router-ports <portnum>

NOTE: To define multiple router ports on a switch, enter the port numbers, separated by blanks. You can enter up to eight router ports in a single command line. To enter more than eight ports, enter the **server router-port** command again with the additional ports.

Defining the Firewalls and Adding them to the Firewall Group

When FWLB is enabled, all the ServerIron ports are in firewall group 2 by default. However, you need to add an entry for each firewall. To add an entry for a firewall, specify the firewall name and IP address. You can specify a name up to 32 characters long. After you add the firewall entries, add the firewalls to the firewall group.

To define the firewalls shown in Figure 7.2 on page 7-9, use the following method.

USING THE CLI

To define the firewalls using the CLI, enter the following commands:

Commands for Active ServerIron A (External Active)

```
SI-ActiveA(config)# server fw-name fw1 192.168.1.2
SI-ActiveA(config-rs-fw1)# exit
SI-ActiveA(config)# server fw-name fw2 192.168.1.3
SI-ActiveA(config-rs-fw2)# exit
SI-ActiveA(config)# server fw-group 2
SI-ActiveA(config-tc-2)# fw-name fw1
SI-ActiveA(config-tc-2)# fw-name fw2
```

Commands for Standby ServerIron A (External Standby)

```
SI-StandbyA(config)# server fw-name fw1 192.168.2.2
SI-StandbyA(config-rs-fw1)# exit
SI-StandbyA(config)# server fw-name fw2 192.168.2.3
SI-StandbyA(config-rs-fw2)# exit
SI-StandbyA(config)# fw-group 2
SI-StandbyA(config-tc-2)# fw-name fw1
SI-StandbyA(config-tc-2)# fw-name fw2
```

Commands for Active ServerIron B (Internal Active)

```
SI-ActiveB(config)# server fw-name fw1 4.4.4.10
SI-ActiveB(config-rs-fw1)# exit
SI-ActiveB(config)# server fw-name fw2 4.4.4.11
SI-ActiveB(config-rs-fw2)# exit
SI-ActiveB(config)# server fw-group 2
SI-ActiveB(config-tc-2)# fw-name fw1
SI-ActiveB(config-tc-2)# fw-name fw2
```

Commands for Standby ServerIron B (Internal Standby)

```
SI-StandbyB(config)# server fw-name fw1 3.3.3.10
SI-StandbyB(config-rs-fw1)# exit
SI-StandbyB(config)# server fw-name fw2 3.3.3.11
SI-StandbyB(config-rs-fw2)# exit
SI-StandbyB(config)# server fw-group 2
SI-StandbyB(config-tc-2)# fw-name fw1
SI-StandbyB(config-tc-2)# fw-name fw2
```

Command Syntax

Syntax: [no] server fw-name <string> <ip-addr>

NOTE: When you add a firewall name, the CLI level changes to the Firewall level. This level is used when you are configuring stateful FWLB.

Syntax: server fw-group 2

This command changes the CLI to firewall group configuration level. The firewall group number is 2. Only one firewall group is supported.

Syntax: [no] fw-name <string>

Adds a configured firewall to the firewall group.

Configuring Paths and Adding Static MAC Entries for Layer 3 Firewalls

A path is configuration information the ServerIron uses to ensure that a given source and destination IP pair is always authenticated by the same Layer 3 firewall.

Each path consists of the following parameters:

- The path ID – A number that identifies the path. In basic FWLB configurations, the paths go from one ServerIron to the other through the firewalls. The paths go from one ServerIron to the ServerIrons in the other active-standby pair other through the firewalls. A path also goes to the router.
- The ServerIron port – The number of the port that connects the ServerIron to the firewall.
- The other ServerIron's or Layer 2 switch's IP address – The management address of the ServerIron or Layer 2 switch on the other side of the firewall. The ServerIron on the private network side and the other ServerIron or Layer 2 switch are the end points of the data path through the firewall.
- The next-hop IP address – The IP address of the firewall interface connected to this ServerIron.

For each type of firewall (Layer 3 synchronous and asynchronous, with or without NAT), you must configure paths between the ServerIrons through the firewalls.

In addition to configuring the paths, you need to create a static MAC entry for each firewall interface attached to the ServerIron.

NOTE: FWLB paths must be fully meshed. When you configure a FWLB path on a ServerIron, make sure you also configure a reciprocal path on the ServerIron attached to the other end of the firewalls. For example, if you configure four paths to four separate firewalls, make sure you configure four paths on the other ServerIron.

NOTE: The static MAC entries are required. You must add a static MAC entry for each firewall interface with the ServerIron.

To configure a path and add static MAC entries, use one of the following methods.

USING THE CLI

To configure the paths and static MAC entries for the configuration shown in Figure 7.2 on page 7-9, enter the following commands. Enter the first group of commands on ServerIron A. Enter the second group of commands on ServerIron B.

Commands for Active ServerIron A (External Active)

```
SI-ActiveA(config)# server fw-group 2
SI-ActiveA(config-tc-2)# fwall-info 1 1 3.3.3.20 192.168.1.2
SI-ActiveA(config-tc-2)# fwall-info 2 2 3.3.3.20 192.168.1.3
SI-ActiveA(config-tc-2)# fwall-info 3 1 4.4.4.20 192.168.1.2
SI-ActiveA(config-tc-2)# fwall-info 4 2 4.4.4.20 192.168.1.3
SI-ActiveA(config-tc-2)# fwall-info 5 8 192.168.1.1 192.168.1.1
SI-ActiveA(config-tc-2)# exit
SI-ActiveA(config)# static-mac-address abcd.4321.2498 ethernet 1 high-priority
router-type
SI-ActiveA(config)# static-mac-address abcd.4321.a53c ethernet 2 high-priority
router-type
```


Commands for Standby ServerIron A (External Standby)

```
SI-StandbyA(config)# server fw-group 2
SI-StandbyA(config-tc-2)# fwall-info 1 1 3.3.3.20 192.168.2.2
SI-StandbyA(config-tc-2)# fwall-info 2 2 3.3.3.20 192.168.2.3
SI-StandbyA(config-tc-2)# fwall-info 3 1 4.4.4.20 192.168.2.2
SI-StandbyA(config-tc-2)# fwall-info 4 2 4.4.4.20 192.168.2.3
SI-StandbyA(config-tc-2)# fwall-info 5 8 192.168.2.1 192.168.2.1
SI-StandbyA(config-tc-2)# exit
SI-StandbyA(config)# static-mac-address abcd.4321.a53d ethernet 2 high-priority
router-type
SI-StandbyA(config)# static-mac-address abcd.4321.2499 ethernet 1 high-priority
router-type
```

Commands for Active ServerIron B (Internal Active)

```
SI-ActiveB(config)# server fw-group 2
SI-ActiveB(config-tc-2)# fwall-info 1 1 192.168.2.10 4.4.4.10
SI-ActiveB(config-tc-2)# fwall-info 2 2 192.168.2.10 4.4.4.11
SI-ActiveB(config-tc-2)# fwall-info 3 1 192.168.1.10 4.4.4.10
SI-ActiveB(config-tc-2)# fwall-info 4 2 192.168.1.10 4.4.4.11
SI-ActiveB(config-tc-2)# fwall-info 5 8 4.4.4.30 4.4.4.30
SI-ActiveB(config-tc-2)# exit
SI-ActiveB(config)# static-mac-address abcd.4321.249b ethernet 1 high-priority
router-type
SI-ActiveB(config)# static-mac-address abcd.4321.a53f ethernet 2 high-priority
router-type
```

Commands for Standby ServerIron B (Internal Standby)

```
SI-StandbyB(config)# server fw-group 2
SI-StandbyB(config-tc-2)# fwall-info 1 1 192.168.1.10 3.3.3.10
SI-StandbyB(config-tc-2)# fwall-info 2 2 192.168.1.10 3.3.3.11
SI-StandbyB(config-tc-2)# fwall-info 3 1 192.168.2.10 3.3.3.10
SI-StandbyB(config-tc-2)# fwall-info 4 2 192.168.2.10 3.3.3.11
SI-StandbyB(config-tc-2)# fwall-info 5 8 3.3.3.30 3.3.3.30
SI-StandbyB(config-tc-2)# exit
SI-StandbyB(config)# static-mac-address abcd.4321.a53e ethernet 2 high-priority
router-type
SI-StandbyB(config)# static-mac-address abcd.4321.249a ethernet 1 high-priority
router-type
```

Command Syntax

Syntax: server fw-group 2

Syntax: [no] fwall-info <path-num> <portnum> <other-ServerIron-ip> <next-hop-ip>

The syntax for adding static MAC entries differs depending on whether you are using a stackable or chassis ServerIron.

Syntax for chassis devices:

Syntax: [no] static-mac-address <mac-addr> ethernet <portnum> [priority <0-7>] [host-type | router-type]

Syntax for stackable devices:

Syntax: static-mac-address <mac-addr> ethernet <portnum> [to <portnum> ethernet <portnum>] [normal-priority | high-priority] [host-type | router-type | fixed-host]

The priority can be 0 – 7 (0 is lowest and 7 is highest) for chassis devices and either normal-priority or high-priority for stackable devices.

The defaults are **host-type** and **0** or **normal-priority**.

NOTE: The static MAC entries are required. You must add a static MAC entry for each firewall interface with the ServerIron. In addition, you must use the **high-priority** and **router-type** parameters with the **static-mac-address** command. These parameters enable the ServerIron to use the address for FWLB.

NOTE: If you enter the command at the global CONFIG level, the static MAC entry applies to the default port-based VLAN (VLAN 1). If you enter the command at the configuration level for a specific port-based VLAN, the entry applies to that VLAN and not to the default VLAN.

Configuring the ServerIron Priority

If you are configuring the ServerIron for IronClad FWLB, you need to specify the priority for the firewalls within the firewall group. The priority determines which of the partner ServerIrons that are configured together for IronClad FWLB is the default active ServerIron for the firewalls within the group.

You can specify a priority from 0 – 255. The ServerIron with the higher priority is the default active ServerIron for the firewalls within the firewall group.

NOTE: If you specify 0, the CLI removes the priority. When you save the configuration to the startup-config file, the **sym-priority** command is removed. Use this method to remove the priority. You cannot remove the priority using the **no sym-priority** command.

USING THE CLI

To configure a ServerIron to be the default active ServerIron for the firewalls in group 2, enter the following commands.

Commands for Active ServerIron A (External Active)

```
SI-ActiveA(config)# server fw-group 2
SI-ActiveA(config-tc-2)# sym-priority 255
```

Commands for Standby ServerIron A (External Standby)

To configure another ServerIron to not be the default active ServerIron for the firewalls in group 2, enter the following commands:

```
SI-StandbyA(config)# server fw-group 2
SI-StandbyA(config-tc-2)# sym-priority 1
```

Commands for Active ServerIron B (Internal Active)

```
SI-ActiveB(config)# server fw-group 2
SI-ActiveB(config-tc-2)# sym-priority 255
```

Commands for Standby ServerIron B (Internal Standby)

```
SI-StandbyB(config)# server fw-group 2
SI-StandbyB(config-tc-2)# sym-priority 1
```

Command Syntax

Syntax: [no] sym-priority <num>

The priority can be from 0 – 255.

NOTE: If you specify 0, the CLI removes the priority. When you save the configuration to the startup-config file, the **sym-priority** command is removed. Use this method to remove the priority. You cannot remove the priority using the **no sym-priority** command.

Preventing Load Balancing of the NAT Addresses

When you configure ServerIrons for load balancing traffic across NAT firewalls, you must disable load balancing on the NAT addresses themselves. You can use either of the following methods to do so. Each method is equally valid and only one of the methods is required. You need to use one of these methods only on the ServerIron connected to the external network, not the ServerIron on the internal side of the network.

- Configure the NAT addresses as firewall addresses, but do not configure paths for the addresses. (This is shown below in the "Extra Firewall Method" section.)
- Configure IP access policies (filters) to deny load balancing for traffic addressed to the NAT addresses. (This is shown below in the "Access Policy Method" section.)

NOTE: In FWLB configurations, the IP policies do not block traffic altogether. They deny load balancing for the traffic. Thus, the ServerIron does not load balance packets addressed to the NAT addresses, but instead sends the traffic only to the firewall that originally sent the traffic.

USING THE CLI

Use either of the following methods to disable load balancing for the NAT addresses.

Extra Firewall Method

To disable load balancing for the NAT addresses by adding firewalls for the addresses, enter commands such as the following.

NOTE: Do not configure paths for the firewalls.

```
SI-ActiveA(config)# server fw-name fw1NAT 192.168.3.1
SI-ActiveA(config-rs-fw1NAT)# exit
SI-ActiveA(config)# server fw-name fw2NAT 192.168.2.3
SI-ActiveA(config-rs-fw2NAT)# exit
```

Access Policy Method

To disable load balancing for the NAT addresses using IP access policies, enter commands such as the following.

```
SI-ActiveA(config)# ip filter 1 deny any 192.168.3.1 255.255.255.255
SI-ActiveA(config)# ip filter 2 deny any 192.168.3.2 255.255.255.255
SI-ActiveA(config)# ip filter 1024 permit any any
```

The first two commands configure policies to deny load balancing for the two NAT addresses. The third command allows all other traffic to be load balanced.

NOTE: The third policy, which permits all traffic, is required because once you define an access policy, the default action for packets that do not match a policy is to deny them. Thus, if you configure only the first two policies and not the third one, you actually disable load balancing altogether by denying the load balancing for all packets.

Configuration Example for IronClad FWLB with Layer 3 NAT Firewalls

This section shows the CLI commands for implementing the configuration shown in Figure 7.2 on page 7-9. The only additional step required is to ensure that the ServerIron connected to the external network does not load balance return traffic to the addresses the firewalls use for NAT. For example, ServerIron A in Figure 7.2 on page 7-9 must be configured so that it does not load balance return traffic to 192.168.2.10/24 or 192.168.2.3/24.

To prevent the ServerIron from load balancing the NAT addresses, you can use either of the following methods. Each method is equally valid and only one of the methods is required. You need to use one of these methods only on the ServerIron connected to the external network, not the ServerIron on the internal side of the network.

- Configure the NAT addresses as firewall addresses, but do not configure paths for the addresses.
- Configure IP access policies (filters) to deny load balancing for traffic addressed to the NAT addresses.

NOTE: In FWLB configurations, the IP policies do not block traffic altogether. They deny load balancing for the traffic. Thus, the ServerIron does not load balance packets addressed to the NAT addresses, but instead sends the traffic only to the firewall that originally sent the traffic.

Commands on Active ServerIron A (External Active)

```
SI-ActiveA(config)# ip address 192.168.1.10/24
SI-ActiveA(config)# ip default-gateway 192.168.1.2
```

The commands above add a management IP address and default gateway address to the ServerIron. For the configuration in this example, the ServerIron needs to be in only one sub-net, so additional IP addresses are not added. However, the IP address must be in the same sub-net as the ServerIron's interface to the Layer 3 firewalls.

```
SI-ActiveA(config)# vlan 10 by port
SI-ActiveA(config-vlan-10)# untagged 5 to 6
SI-ActiveA(config-vlan-10)# exit
```

The commands above configure the ports for the connection to the standby ServerIron in a separate port-based VLAN. This is required.

```
SI-ActiveA(config)# trunk switch ethernet 5 to 6
```

The **trunk** command creates a trunk group for the ports that connect this ServerIron to its partner. Using a trunk group for the link between the active and standby ServerIrons is not required, but using a trunk group adds an additional level of redundancy for enhanced availability. If one of the ports in a trunk group goes down, the link remains intact as long as the other port remains up. Since the trunk group is between two ServerIron switches, make sure you configure a switch trunk group, not a server trunk group.

```
SI-ActiveA(config)# server router-port 8
```

The **server router-port** command identifies the port that connects this ServerIron to the router connected to the other ServerIron in the active-standby pair.

```
SI-ActiveA(config)# server fw-port 5
```

The **server fw-port** command identifies the port that connects this ServerIron to its partner. If you configure a trunk group for the link between the two partners, specify the first port (the primary port for the group) in the trunk group. On the 8-port, 16-port, and 24-port ServerIrons, you can configure a trunk group with two or four members and the lead ports are the odd-numbered ports.

```
SI-ActiveA(config)# server fw-name fw1 192.168.1.2
SI-ActiveA(config-rs-fw1)# exit
SI-ActiveA(config)# server fw-name fw2 192.168.1.3
SI-ActiveA(config-rs-fw2)# exit
```

The **server fw-name** commands add the firewalls to the ServerIron. In the commands above, "fw1" and "fw2" are the firewall names. These names are specific to the ServerIron and do not need to correspond to any name parameters on the firewalls themselves. The IP addresses are the addresses of the firewall interfaces with the ServerIron.

The following commands add firewall entries for the hidden NAT addresses. These entries prevent the ServerIron from load balancing the firewall traffic to these addresses. The ServerIron forwards a return packet addressed to one of these firewalls directly to the firewall that sent it, instead of using the hash mechanism to select a path for the traffic.

```
ServerIron-A(config)# server fw-name fw3NAT 192.168.2.10
ServerIron-A(config-rs-fw3NAT)# exit
ServerIron-A(config)# server fw-name fw4NAT 192.168.2.3
```



```
ServerIron-A(config-rs-fw4NAT)# exit
```

The following commands configure the firewall group. The **server fw-group 2** command changes the focus of the CLI to firewall group 2.

The **sym-priority** command specifies the priority of this ServerIron with respect to the other ServerIron for the firewalls in the firewall group. The priority can be from 0 – 255. The ServerIron with the higher priority is the default active ServerIron for the firewalls within the group.

NOTE: If you specify 0, the CLI removes the priority. When you save the configuration to the startup-config file, the **sym-priority** command is removed. Use this method to remove the priority. You cannot remove the priority using the **no sym-priority** command.

The **fw-name** <firewall-name> command adds the firewalls to the firewall group. Notice that the firewall entries for the hidden NAT addresses are not added.

```
SI-ActiveA(config)# server fw-group 2
SI-ActiveA(config-tc-2)# sym-priority 255
SI-ActiveA(config-tc-2)# fw-name fw1
SI-ActiveA(config-tc-2)# fw-name fw2
```

The **fwall-info** commands add the paths between this ServerIron and the other ServerIrons through the firewalls. The paths enhance performance by ensuring that a given traffic flow (source and destination IP addresses) always travels through the same firewall. In configurations that use asynchronous firewalls, the paths enhance performance by eliminating excess authentications. In this configuration, each ServerIron has two paths to each of the two firewalls. The fifth path goes to the router.

The paths are required, even if the firewalls are synchronized.

The first parameter with each command is a path ID. The second parameter is the port number of the ServerIron port that connects the ServerIron to the firewall in the path.

The third parameter is the IP address of the ServerIron at the other end of the path or, for paths to routers, the IP address of the router's interface with the ServerIron. Note that each ServerIron has a path to each of the ServerIrons in the other pair, but does not have a path to its own standby pair.

The fourth parameter is the IP address of the firewall or router interface with this ServerIron. Notice that the ServerIron has two paths for each firewall. One of the paths goes to the active ServerIron in the other pair. The other path goes to the standby ServerIron in the pair. In the case of the path to the router, the third and forth parameters have the same value.

```
SI-ActiveA(config-tc-2)# fwall-info 1 1 3.3.3.20 192.168.1.2
SI-ActiveA(config-tc-2)# fwall-info 2 2 3.3.3.20 192.168.1.3
SI-ActiveA(config-tc-2)# fwall-info 3 1 4.4.4.20 192.168.1.2
SI-ActiveA(config-tc-2)# fwall-info 4 2 4.4.4.20 192.168.1.3
SI-ActiveA(config-tc-2)# fwall-info 5 8 192.168.1.1 192.168.1.1
SI-ActiveA(config-tc-2)# exit
```

The commands below add static entries to the ServerIron's MAC table for the firewall interfaces. The high-priority and router-type parameters are required for FWLB.

```
SI-ActiveA(config)# vlan 1
SI-ActiveA(config-vlan-1)# static-mac-address abcd.4321.2498 ethernet 1 high-
priority router-type
SI-ActiveA(config-vlan-1)# static-mac-address abcd.4321.a53c ethernet 2 high-
priority router-type
SI-ActiveA(config-vlan-1)# exit
```

NOTE: If you enter the command at the global CONFIG level, the static MAC entry applies to the default port-based VLAN (VLAN 1). If you enter the command at the configuration level for a specific port-based VLAN, the entry applies to that VLAN and not to the default VLAN.

The commands below globally enable firewall balancing. The "0" parameter is required and enables the ServerIron to provide FWLB for all packets of the specified type (TCP or UDP). The **write memory** command saves the configuration changes made by all these commands to the ServerIron's startup-config file.


```
SI-ActiveA(config)# ip policy 1 fw tcp 0 global
SI-ActiveA(config)# ip policy 2 fw udp 0 global
SI-ActiveA(config)# write memory
```

Alternative Configuration for Active ServerIron A

The example above configures FWLB for NAT firewalls by adding firewall definitions for the IP addresses the NAT service on the firewalls uses for traffic sent from a client inside the firewalls to a destination outside the firewalls.

Alternatively, you can configure IP access policies that deny load balancing for the NAT addresses. For the example in Figure 7.2 on page 7-9, you would enter the following commands:

```
ServerIron-A(config)# ip filter 1 deny any 192.168.2.3 255.255.255.255
ServerIron-A(config)# ip filter 2 deny any 192.168.3.2 255.255.255.255
ServerIron-A(config)# ip filter 1024 permit any
```

The first two commands configure policies to deny load balancing for the two NAT addresses. The third command allows all other traffic to be load balanced.

NOTE: The third policy, which permits all traffic, is required because once you define an access policy, the default action for packets that do not match a policy is to deny them. Thus, if you configure only the first two policies and not the third one, you actually disable load balancing altogether by denying the load balancing for all packets.

The other commands are the same as in the previous section.

Commands on Standby ServerIron A (External Standby)

```
SI-StandbyA(config)# ip address 192.168.2.10/24
SI-StandbyA(config)# ip default-gateway 192.168.2.2
SI-StandbyA(config)# vlan 10 by port
SI-StandbyA(config-vlan-10)# untagged 5 to 6
SI-StandbyA(config-vlan-10)# exit
SI-StandbyA(config)# trunk switch ethernet 5 to 6
SI-StandbyA(config)# server router-port 8
SI-StandbyA(config)# server fw-port 5
SI-StandbyA(config)# server fw-name fw2-1 192.168.2.2
SI-StandbyA(config-rs-fw2-1)# exit
SI-StandbyA(config)# server fw-name fw2-2 192.168.2.3
SI-StandbyA(config-rs-fw2-2)# exit
SI-StandbyA(config)# server fw-group 2
SI-StandbyA(config-tc-2)# sym-priority 1
SI-StandbyA(config-tc-2)# fw-name fw1
SI-StandbyA(config-tc-2)# fw-name fw2
SI-StandbyA(config-tc-2)# fwall-info 1 1 3.3.3.20 192.168.2.2
SI-StandbyA(config-tc-2)# fwall-info 2 2 3.3.3.20 192.168.2.3
SI-StandbyA(config-tc-2)# fwall-info 3 1 4.4.4.20 192.168.2.2
SI-StandbyA(config-tc-2)# fwall-info 4 2 4.4.4.20 192.168.2.3
SI-StandbyA(config-tc-2)# fwall-info 5 8 192.168.2.1 192.168.2.1
SI-StandbyA(config-tc-2)# exit
SI-StandbyA(config)# vlan 1
SI-StandbyA(config-vlan-1)# static-mac-address abcd.4321.a53d ethernet 2 high-
priority router-type
SI-StandbyA(config-vlan-1)# static-mac-address abcd.4321.2499 ethernet 1 high-
priority router-type
SI-StandbyA(config-vlan-1)# exit
SI-StandbyA(config)# ip policy 1 fw tcp 0 global
SI-StandbyA(config)# ip policy 2 fw udp 0 global
SI-StandbyA(config)# write memory
```


Alternative Configuration for Standby ServerIron A

The example above configures FWLB for NAT firewalls by adding firewall definitions for the IP addresses the NAT service on the firewalls uses for traffic sent from a client inside the firewalls to a destination outside the firewalls.

Alternatively, you can configure IP access policies that deny load balancing for the NAT addresses. For the example in Figure 7.2 on page 7-9, you would enter the following commands:

```
SI-StandbyA(config)# ip filter 1 deny any 192.168.2.3 255.255.255.255
SI-StandbyA(config)# ip filter 2 deny any 192.168.3.2 255.255.255.255
SI-StandbyA(config)# ip filter 1024 permit any any
```

The first two commands configure policies to deny load balancing for the two NAT addresses. The third command allows all other traffic to be load balanced.

NOTE: The third policy, which permits all traffic, is required because once you define an access policy, the default action for packets that do not match a policy is to deny them. Thus, if you configure only the first two policies and not the third one, you actually disable load balancing altogether by denying the load balancing for all packets.

The other commands are the same as in the previous section.

Commands on Active ServerIron B (Internal Active)

```
SI-ActiveB(config)# ip address 3.3.3.20/24
SI-ActiveB(config)# ip default-gateway 4.4.4.11
SI-ActiveB(config)# vlan 10 by port
SI-ActiveB(config-vlan-10)# untagged 5 to 6
SI-ActiveB(config-vlan-10)# exit
SI-ActiveB(config)# trunk switch ethernet 5 to 6
SI-ActiveB(config)# server router-port 8
SI-ActiveB(config)# server fw-port 5
SI-ActiveB(config)# server fw-name fw2-1 4.4.4.10
SI-ActiveB(config-rs-fw2-1)# exit
SI-ActiveB(config)# server fw-name fw2-2 4.4.4.11
SI-ActiveB(config-rs-fw2-2)# exit
SI-ActiveB(config)# server fw-group 2
SI-ActiveB(config-tc-2)# sym-priority 255
SI-ActiveB(config-tc-2)# fw-name fw2-1
SI-ActiveB(config-tc-2)# fw-name fw2-2
SI-ActiveB(config-tc-2)# fwall-info 1 1 192.168.2.10 4.4.4.10
SI-ActiveB(config-tc-2)# fwall-info 2 2 192.168.2.10 4.4.4.11
SI-ActiveB(config-tc-2)# fwall-info 3 1 192.168.1.10 4.4.4.10
SI-ActiveB(config-tc-2)# fwall-info 4 2 192.168.1.10 4.4.4.11
SI-ActiveB(config-tc-2)# fwall-info 5 8 4.4.4.30 4.4.4.30
SI-ActiveB(config-tc-2)# exit
SI-ActiveB(config)# vlan 1
SI-ActiveB(config-vlan-1)# static-mac-address abcd.4321.249b ethernet 1 high-
priority router-type
SI-ActiveB(config-vlan-1)# static-mac-address abcd.4321.a53f ethernet 2 high-
priority router-type
SI-ActiveB(config-vlan-1)# exit
SI-ActiveB(config)# ip policy 1 fw tcp 0 global
SI-ActiveB(config)# ip policy 2 fw udp 0 global
SI-ActiveB(config)# write memory
```


Chapter 8
Configuring FWLB and SLB

NOTE: This chapter shows basic FWLB configurations with Layer 3 firewalls. Currently, these are the configurations supported by the ServerIron. If you need to perform concurrent SLB and FWLB in a different type of FWLB configuration, contact Brocade Communications Systems.

You can configure the ServerIron to concurrently perform FWLB and SLB at the same time. The software supports the following configurations:

- SLB-to-FWLB – The ServerIron on the Internet side of the firewalls performs FWLB for traffic directed toward real servers connected to the ServerIron on the private side of the firewalls. In this configuration, all the SLB configuration (virtual IP address, real server, and port bindings) resides on the Internet ServerIron. The real servers are configured as remote servers. In addition, the SLB-to-FWLB feature is enabled on the Internet ServerIron. The internal ServerIron is configured for FWLB but requires no additional configuration.
- FWLB-to-SLB – The internal ServerIron (the one on the private side of the firewalls) contains all the SLB configuration information. In this configuration, the FWLB-to-SLB feature is enabled on this ServerIron rather than the Internet ServerIron. This configuration enables the internal ServerIron to learn the firewall from which a client request is received and send the server reply back through the same firewall.

Your choice of implementation depends on the ServerIron you want to use for the SLB configuration. Use SLB-to-FWLB is you want to place the SLB configuration on the Internet ServerIron. Use FWLB-to-SLB is you want to place the SLB configuration on the internal ServerIron.

NOTE: You must use hash-based FWLB (the default) if you use either of these features. The ServerIron does not support stateful FWLB with these features.

NOTE: On the ServerIronXL, you must use the default VLAN (normally VLAN 1) for the FWLB configuration.

Figure 8.1 shows an example of an SLB-to-FWLB configuration.

Figure 8.1 SLB-to-FWLB configuration

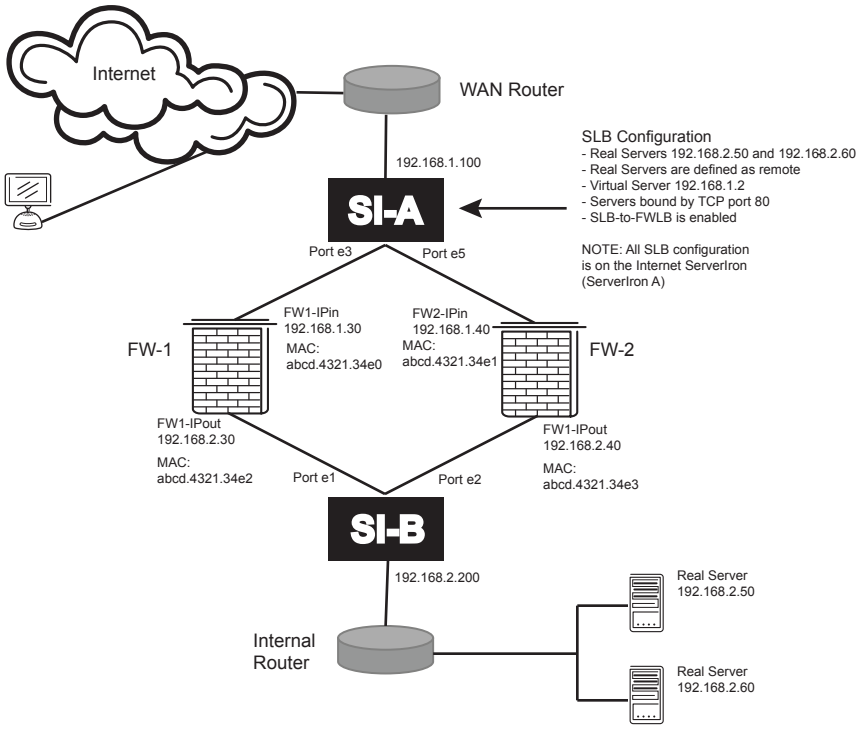

Notice that all the SLB configuration is on the Internet ServerIron (ServerIron A).

Figure 8.2 shows an example of an SLB-to-FWLB configuration.

Figure 8.2 FWLB-to-SLB configuration

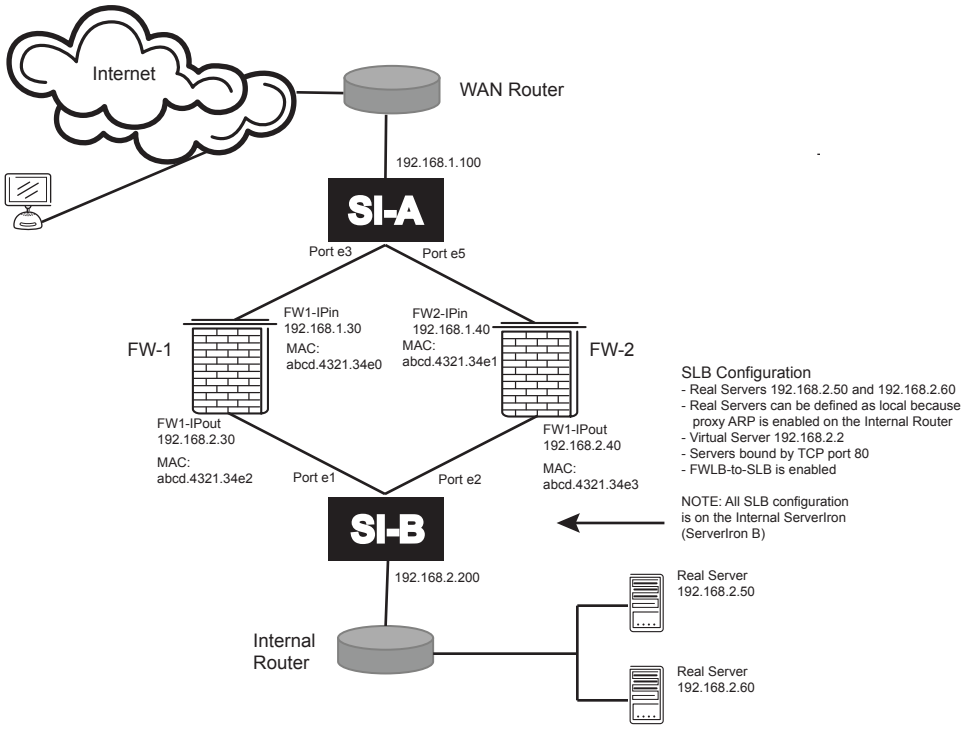

For FWLB-to-SLB, all the SLB configuration information is on the internal ServerIron (ServerIron B).

Configuring SLB-to-FWLB

To configure SLB-to-FWLB in a basic FWLB configuration for Layer 3 firewalls, such as the one shown in Figure 8.1, perform the following tasks.

- **Configure SLB parameters on the Internet ServerIron**
  - Configure the real servers
  - Configure the virtual server
  - Bind the real servers to the virtual server
  - Enable the SLB-to-FWLB feature
- **Configure global FWLB parameters**
  - Globally enable FWLB
- **Configure firewall parameters**
  - Define the firewalls and add them to the firewall group
- **Configure firewall group parameters**
  - Configure the paths and add static MAC entries for the firewall interfaces with the ServerIron

NOTE: On the ServerIronXL, you must use the default VLAN (normally VLAN 1) for the FWLB configuration.

The tasks under the first item (Configure SLB parameters on the Internet ServerIron) are described in the following sections. The remaining tasks are identical to the tasks for configuring basic FWLB for Layer 3 firewalls. For more information about these tasks, see “Configuring Basic Layer 3 FWLB” on page 4-1.

Configuring the SLB Parameters

In an SLB-to-FWLB configuration, all SLB configuration takes place on the Internet ServerIron. The ServerIron on the private side of the firewalls does not contain any SLB configuration information. This section describes how to configure the Internet ServerIron to provide SLB for the real servers and virtual server shown in Figure 8.1 on page 8-2.

NOTE: This section describes basic SLB configuration tasks. For advanced configuration features, see the *ServerIron TrafficWorks Server Load Balancing Guide*.

Configuring the Real Servers

To configure the real servers shown in Figure 8.1 on page 8-2, enter the following commands on the Internet ServerIron (ServerIron A).

NOTE: In SLB-to-FWLB configurations, you must define the real servers as remote servers.

USING THE CLI

```
ServerIronA(config)# server remote-name RS1 192.168.2.50
ServerIronA(config-rs-RS1)# port http
ServerIronA(config-rs-RS1)# exit
ServerIronA(config)# server remote-name RS2 192.168.2.60
ServerIronA(config-rs-RS2)# port http
ServerIronA(config-rs-RS2)# exit
```

The **server remote-name** command adds a real server. The port command enables a TCP or UDP port on the server. In this case, the **port http** command enables TCP port 80 (HTTP).

NOTE: If you use the **server real-name** command instead of the **server remote-name** command, the real servers are added as local servers. Use the **server remote-name** command. You must add them as remote servers for SLB-to-FWLB.

Syntax: [no] server remote-name <text> <ip-addr>

Syntax: [no] port <port> [disable | enable]

Syntax: [no] port <port> [keepalive]

For information about the optional parameters with the **port** command, see the "Health Check" chapter in the *ServerIron TrafficWorks Server Load Balancing Guide*.

Configuring the Virtual Server

To configure the virtual server shown in Figure 8.1 on page 8-2, enter the following command on the Internet ServerIron (ServerIron A).

USING THE CLI

```
ServerIronA(config)# server virtual-name www.brocade.com 192.168.1.2
ServerIronA(config-vs-www.brocade.com)# port http
```

The **server virtual-name** command adds the virtual server. The **port** command enables a TCP or UDP port on the server.

Syntax: [no] server virtual-name <text> [<ip-addr>]

Binding the Real Server to the Virtual Server

To bind the real servers to the virtual server, enter the following commands on the Internet ServerIron (ServerIron A). Notice that the port binding takes place on the Virtual Server configuration level.

USING THE CLI

```
ServerIronA(config)# server virtual www.brocade.com
ServerIronA(config-vs-www.brocade.com)# bind http RS1 http
ServerIronA(config-vs-www.brocade.com)# bind http RS2 http
ServerIronA(config-vs-www.brocade.com)# exit
```

Syntax: [no] bind <port> <real server name> <port>

Enabling SLB-to-FWLB

To enable SLB-to-FWLB, enter the following command on the Internet ServerIron (ServerIron A).

```
ServerIronA(config)# server slb-fw
```

Syntax: [no] server slb-fw

Configuration Example for SLB-to-FWLB

The following sections show all the ServerIron commands you would enter on each ServerIron to implement the SLB-to-FWLB configuration shown in Figure 8.1 on page 8-2.

Commands on ServerIron A (External)

Enter the following commands to configure SLB. In an SLB-to-FWLB configuration, all SLB configuration takes place on the Internet ServerIron (ServerIron A, the External ServerIron, in this example).

The following commands change the ServerIron's host name to "ServerIronA", configure the ServerIron's management IP address, and specify the default gateway.

```
ServerIron(config)# hostname ServerIronA
ServerIronA(config)# ip address 192.168.1.100 255.255.255.0
ServerIronA(config)# ip default-gateway 192.168.1.1
```

The following commands configure the real servers. Notice that the servers are configured as remote servers. This is required for SLB-to-FWLB.

```
ServerIronA(config)# server remote-name RS1 192.168.2.50
ServerIronA(config-rs-RS1)# port http
ServerIronA(config-rs-RS1)# exit
ServerIronA(config)# server remote-name RS2 192.168.2.60
ServerIronA(config-rs-RS2)# port http
ServerIronA(config-rs-RS2)# exit
```

The following commands configure the virtual server and bind it to the real servers with TCP port 80 (HTTP).

```
ServerIronA(config)# server virtual-name www.brocade.com 192.168.1.2
ServerIronA(config-vs-www.brocade.com)# port http
ServerIronA(config)# server virtual www.brocade.com
ServerIronA(config-vs-www.brocade.com)# bind http RS1 http
ServerIronA(config-vs-www.brocade.com)# bind http RS2 http
```

Enter the following command to enable SLB-to-FWLB.

NOTE: This command applies only to the ServerIron that contains the SLB configuration. Do not enter this command on the internal ServerIron (ServerIronB).

```
ServerIronA(config)# server slb-fw
```

The the following commands to add two firewalls, FW1-IPin and FW2-IPin.


```
ServerIronA(config)# server fw-name FW1-IPin 192.168.1.30
ServerIronA(config-rs-FW1-IPin)# exit
ServerIronA(config)# server fw-name FW2-IPin 192.168.1.40
ServerIronA(config-rs-FW2-IPin)# exit
```

The following commands configure parameters for firewall group 2. The **fwall-info** commands configure the paths for the firewall traffic. Each path consists of a path ID, the ServerIron port attached to the firewall, the IP address of the ServerIron at the other end of the path, and the next-hop IP address (usually the firewall interface connected to this ServerIron). Make sure you configure reciprocal paths on the other ServerIron, as shown in the section containing the CLI commands for ServerIron B.

NOTE: Path information is required even if the firewalls are synchronized.

The **fw-name** <firewall-name> command adds the firewalls to the firewall group.

```
ServerIronA(config)# server fw-group 2
ServerIronA(config-tc-2)# fw-name FW1-IPin
ServerIronA(config-tc-2)# fw-name FW2-IPin
ServerIronA(config-tc-2)# fwall-info 1 3 192.168.2.200 192.168.1.30
ServerIronA(config-tc-2)# fwall-info 2 5 192.168.2.200 192.168.1.40
ServerIronA(config-tc-2)# exit
```

The following commands add static MAC entries for the MAC addresses of the firewall interfaces connected to the ServerIron. Notice that the QoS priority is configured as **high-priority** and the **router-type** parameter is specified. These parameters are required. You must specify **high-priority** and **router-type**.

NOTE: To ensure proper operation, always configure the path IDs so that the IDs consistently range from lowest path ID to highest path ID for the firewalls. For example, in Figure 8.1 on page 8-2, the path IDs should range from lowest to highest beginning with the firewall interface at the upper left of the figure.

To ensure smooth operation, you might want to depict your firewalls in a vertical hierarchy as in Figure 8.1 on page 8-2, label the interfaces with their IP addresses, then configure the paths so that the path IDs to the interfaces range from lowest to highest path ID starting from the uppermost firewall interface.

```
ServerIronA(config)# static-mac-address abcd.4321.34e0 ethernet 3 high-priority
router-type
ServerIronA(config)# static-mac-address abcd.4321.34e1 ethernet 5 high-priority
router-type
```

The following commands configure global policies to enable FWLB. Global or local policies are required for FWLB. The first **ip policy** command in this example configures the ServerIron to perform FWLB for all TCP traffic. The value “0” is equivalent to “any” and means the ServerIron should perform FWLB for all TCP traffic. The second **ip policy** command enables FWLB for all UDP traffic.

```
ServerIronA(config)# ip policy 1 fw tcp 0 global
ServerIronA(config)# ip policy 2 fw udp 0 global
ServerIronA(config)# write memory
```

Commands on ServerIron B (Internal)

Enter the following commands to configure FWLB on ServerIron B. Notice that the **fwall-info** commands configure paths that are reciprocal to the paths configured on ServerIron A. Path 1 on each ServerIron goes through one of the firewalls while path 2 goes through the other firewall.

```
ServerIronB(config)# server fw-name FW1-IPout 192.168.2.30
ServerIronB(config-rs-FW1-IPout)# exit
ServerIronB(config)# server fw-name FW2-IPout 192.168.2.40
ServerIronB(config-rs-FW2-IPout)# exit
ServerIronB(config)# server fw-group 2
ServerIronB(config-tc-2)# fw-name FW1-IPout
ServerIronB(config-tc-2)# fw-name FW2-IPout
ServerIronB(config-tc-2)# fwall-info 1 1 192.168.1.100 192.168.2.30
```



```
ServerIronB(config-tc-2)# fwall-info 2 2 192.168.1.100 192.168.2.40
ServerIronB(config-tc-2)# exit
ServerIronB(config)# static-mac-address abcd.4321.34e2 ethernet 1 high-priority
router-type
ServerIronB(config)# static-mac-address abcd.4321.34e3 ethernet 2 high-priority
router-type
ServerIronB(config)# ip policy 1 fw tcp 0 global
ServerIronB(config)# ip policy 2 fw udp 0 global
ServerIronB(config)# write memory
```

Configuring FWLB-to-SLB

Configuration for FWLB-to-SLB is similar to configuration for SLB-to-FWLB, with the following differences:

- SLB configuration information resides on the internal ServerIron, not on the Internet ServerIron.
- The FWLB-to-SLB feature is enabled on the internal ServerIron.
- If Proxy ARP is enabled on the internal router, you can define the real servers as local servers instead of remote servers. However, if Proxy ARP is not enabled on the internal router, the real servers must be remote servers.

To configure FWLB-to-SLB in a basic FWLB configuration for Layer 3 firewalls, such as the one shown in Figure 8.2 on page 8-3, perform the following tasks.

- **Configure SLB parameters on the internal ServerIron**
  - Configure the real servers
  - Configure the virtual server
  - Bind the real servers to the virtual server
  - Enable the FWLB-to-SLB feature
- **Configure global FWLB parameters**
  - Globally enable FWLB
- **Configure firewall parameters**
  - Define the firewalls and add them to the firewall group
- **Configure firewall group parameters**
  - Configure the paths and add static MAC entries for the firewall interfaces with the ServerIron

NOTE: On the ServerIronXL, you must use the default VLAN (normally VLAN 1) for the FWLB configuration.

The tasks under the first item (Configure SLB parameters on the internal ServerIron) are described in the following sections. The remaining tasks are identical to the tasks for configuring basic FWLB for Layer 3 firewalls. For more information about these tasks, see “Configuring Basic Layer 3 FWLB” on page 4-1.

Configuring the SLB Parameters

In an FWLB-to-SLB configuration, all SLB configuration takes place on the internal ServerIron. The ServerIron on the Internet side of the firewalls does not contain any SLB configuration information. This section describes how to configure the internal ServerIron to provide SLB for the real servers and virtual server shown in Figure 8.2 on page 8-3.

NOTE: This section describes basic SLB configuration tasks. For advanced configuration features, see the *ServerIron TrafficWorks Server Load Balancing Guide*.

Configuring the Real Servers

To configure the real servers shown in Figure 8.2 on page 8-3, enter the following commands on the internal ServerIron (ServerIron B).

NOTE: In FWLB-to-SLB configurations, you must define the real servers as remote servers unless Proxy ARP is enabled on the internal router.

USING THE CLI

```
ServerIronB(config)# server real-name RS1 192.168.2.50
ServerIronB(config-rs-RS1)# port http
ServerIronB(config-rs-RS1)# exit
ServerIronB(config)# server real-name RS2 192.168.2.60
ServerIronB(config-rs-RS2)# port http
ServerIronB(config-rs-RS2)# exit
```

The **server real-name** command adds a real server. The port command enables a TCP or UDP port on the server. In this case, the **port http** command enables TCP port 80 (HTTP).

NOTE: You can use the **server real-name** command if Proxy ARP is enabled on the internal router. Otherwise, you must use the **server remote-name** command to add the real servers instead of the **server real-name** command.

Syntax: [no] server remote-name <text> <ip-addr>

Syntax: [no] port <port> [disable | enable]

Syntax: [no] port <port> [keepalive]

For information about the optional parameters with the **port** command, see the "Health Checks" chapter in the *ServerIron TrafficWorks Server Load Balancing Guide*.

Configuring the Virtual Server

To configure the virtual server shown in Figure 8.2 on page 8-3, enter the following command on the internal ServerIron (ServerIron B).

USING THE CLI

```
ServerIronB(config)# server virtual-name www.brocade.com 192.168.1.2
ServerIronB(config-vs-www.brocade.com)# port http
```

The **server virtual-name** command adds the virtual server. The **port** command enables a TCP or UDP port on the server.

Syntax: [no] server virtual-name <text> [<ip-addr>]

Binding the Real Server to the Virtual Server

To bind the real servers to the virtual server, enter the following commands on the internal ServerIron (ServerIron B). Notice that the port binding takes place on the Virtual Server configuration level.

USING THE CLI

```
ServerIronB(config)# server virtual www.foundrynet.com
ServerIronB(config-vs-www.brocade.com)# bind http RS1 http
ServerIronB(config-vs-www.brocade.com)# bind http RS2 http
ServerIronB(config-vs-www.brocade.com)# exit
```

Syntax: [no] bind <port> <real server name> <port>

Enabling FWLB-to-SLB

To enable FWLB-to-SLB, enter the following command on the internal ServerIron (ServerIron B).


```
ServerIronB(config)# server fw-slb
```

Syntax: [no] server fw-slb

Configuration Example for FWLB-to-SLB

The following sections show all the ServerIron commands you would enter on each ServerIron to implement the FWLB-to-SLB configuration shown in Figure 8.2 on page 8-3.

Commands on ServerIron A (External)

The following commands change the ServerIron's host name to "ServerIronA", configure the ServerIron's management IP address, and specify the default gateway.

```
ServerIron(config)# hostname ServerIronA
ServerIronA(config)# ip address 192.168.1.100 255.255.255.0
ServerIronA(config)# ip default-gateway 192.168.1.1
```

Enter the following commands to add two firewalls, FW1-IPin and FW2-IPin.

```
ServerIronA(config)# server fw-name FW1-IPin 192.168.1.30
ServerIronA(config-rs-FW1-IPin)# exit
ServerIronA(config)# server fw-name FW2-IPin 192.168.1.40
ServerIronA(config-rs-FW2-IPin)# exit
```

The following commands configure parameters for firewall group 2. The **fwall-info** commands configure the paths for the firewall traffic. Each path consists of a path ID, the ServerIron port attached to the firewall, the IP address of the ServerIron at the other end of the path, and the next-hop IP address (usually the firewall interface connected to this ServerIron). Make sure you configure reciprocal paths on the other ServerIron, as shown in the section containing the CLI commands for ServerIron B.

NOTE: Path information is required even if the firewalls are synchronized.

The **fw-name** <firewall-name> command adds the firewalls to the firewall group.

```
ServerIronA(config)# server fw-group 2
ServerIronA(config-tc-2)# fw-name FW1-IPin
ServerIronA(config-tc-2)# fw-name FW2-IPin
ServerIronA(config-tc-2)# fwall-info 1 3 192.168.2.200 192.168.1.30
ServerIronA(config-tc-2)# fwall-info 2 5 192.168.2.200 192.168.1.40
ServerIronA(config-tc-2)# exit
```

The following commands add static MAC entries for the MAC addresses of the firewall interfaces connected to the ServerIron. Notice that the QoS priority is configured as **high-priority** and the **router-type** parameter is specified. These parameters are required. You must specify **high-priority** and **router-type**.

NOTE: To ensure proper operation, always configure the path IDs so that the IDs consistently range from lowest path ID to highest path ID for the firewalls. For example, in Figure 8.2 on page 8-3, the path IDs should range from lowest to highest beginning with the firewall interface at the upper left of the figure.

To ensure smooth operation, you might want to depict your firewalls in a vertical hierarchy as in Figure 8.2 on page 8-3, label the interfaces with their IP addresses, then configure the paths so that the path IDs to the interfaces range from lowest to highest path ID starting from the uppermost firewall interface.

```
ServerIronA(config)# static-mac-address abcd.4321.34e0 ethernet 3 high-priority
router-type
ServerIronA(config)# static-mac-address abcd.4321.34e1 ethernet 5 high-priority
router-type
```

The following commands configure global policies to enable FWLB. Global or local policies are required for FWLB. The first **ip policy** command in this example configures the ServerIron to perform FWLB for all TCP traffic.

The value “0” is equivalent to “any” and means the ServerIron should perform FWLB for all TCP traffic. The second **ip policy** command enables FWLB for all UDP traffic.

```
ServerIronA(config)# ip policy 1 fw tcp 0 global
ServerIronA(config)# ip policy 2 fw udp 0 global
ServerIronA(config)# write memory
```

Commands on ServerIron B (Internal)

Enter the following commands to configure SLB. In an FWLB-to-SLB configuration, all SLB configuration takes place on the internal ServerIron (ServerIron B, the internal ServerIron, in this example).

The following commands change the ServerIron’s host name to “ServerIronB", configure the ServerIron’s management IP address, and specify the default gateway.

```
ServerIron(config)# hostname ServerIronB
ServerIronB(config)# ip address 192.168.2.200 255.255.255.0
ServerIronB(config)# ip default-gateway 192.168.2.1
```

The following commands configure the real servers. Notice that the servers are configured as local servers instead of remote servers. You can configure the real servers as local servers if Proxy ARP is enabled on the internal router.

```
ServerIronB(config)# server real-name RS1 192.168.2.50
ServerIronB(config-rs-RS1)# port http
ServerIronB(config-rs-RS1)# exit
ServerIronB(config)# server real-name RS2 192.168.2.60
ServerIronB(config-rs-RS2)# port http
ServerIronB(config-rs-RS2)# exit
```

The following commands configure the virtual server and bind it to the real servers with TCP port 80 (HTTP).

```
ServerIronB(config)# server virtual-name www.brocade.com 192.168.1.2
ServerIronB(config-vs-www.brocade.com)# port http
ServerIronB(config)# server virtual www.foundrynet.com
ServerIronB(config-vs-www.brocade.com)# bind http RS1 http
ServerIronB(config-vs-www.brocade.com)# bind http RS2 http
```

Enter the following command to enable FWLB-to-SLB.

NOTE: This command applies only to the ServerIron that contains the SLB configuration. Do not enter this command on the Internet ServerIron (ServerIronA).

```
ServerIronB(config)# server fw-slb
```

Enter the following commands to complete the FWLB configuration on this ServerIron. Notice that the **fwall-info** commands configure paths that are reciprocal to the paths configured on ServerIron A. Path 1 on each ServerIron goes through one of the firewalls while path 2 goes through the other firewall.

```
ServerIronB(config)# server fw-name FW1-IPout 192.168.2.30
ServerIronB(config-rs-FW1-IPout)# exit
ServerIronB(config)# server fw-name FW2-IPout 192.168.2.40
ServerIronB(config-rs-FW2-IPout)# exit
ServerIronB(config)# server fw-group 2
ServerIronB(config-tc-2)# fw-name FW1-IPout
ServerIronB(config-tc-2)# fw-name FW2-IPout
ServerIronB(config-tc-2)# fwall-info 1 1 192.168.1.100 192.168.2.30
ServerIronB(config-tc-2)# fwall-info 2 2 192.168.1.100 192.168.2.40
ServerIronB(config-tc-2)# exit
ServerIronB(config)# static-mac-address abcd.4321.34e2 ethernet 1 high-priority
router-type
ServerIronB(config)# static-mac-address abcd.4321.34e3 ethernet 2 high-priority
router-type
ServerIronB(config)# ip policy 1 fw tcp 0 global
```



```
ServerIronB(config)# ip policy 2 fw udp 0 global
ServerIronB(config)# write memory
```

From HA Chapter

Active-Active FWLB – with External SLB (FWLB-to-SLB)

The software supports two types of FWLB with SLB configurations. Your choice of implementation depends on which pair of ServerIrons you want to use for the SLB configuration. Use SLB-to-FWLB is you want to place the SLB configuration on the external ServerIrons. Use FWLB-to-SLB if you want to place the SLB configuration on the internal ServerIrons.

The software supports the following configurations:

- FWLB-to-SLB – The internal ServerIron (the one on the server side or private side of the firewalls) contains all the SLB configuration information. In this configuration, the FWLB-to-SLB feature is enabled on the internal ServerIron rather than the external ServerIron. This configuration enables the internal ServerIron to learn the firewall from which a client request is received and send the server reply back through the same firewall.
- SLB-to-FWLB – The external ServerIron, on the client or external side of the firewalls, performs FWLB for traffic directed toward real servers connected to the ServerIron on the private side of the firewalls. In this configuration, all the SLB configuration (virtual IP address, real server, and port bindings) resides on the external ServerIron. The real servers are configured as remote servers. In addition, the SLB-to-FWLB feature is enabled on the external ServerIron. The internal ServerIron is configured for FWLB but requires no additional configuration.

Figure 8.3 shows an example of an active-active FWLB configuration that also supports SLB. The pair of ServerIrons on the non-secure (external) of the firewalls are connected to clients. The pair of ServerIrons on the secure side of the firewalls are connected to application servers. Both pairs of ServerIrons load balance the traffic to the firewalls and also perform SLB load balancing for application traffic.

Both ServerIrons in each pair actively load balance traffic as well as provide redundancy.

You can configure the network in Figure 8.3 for FWLB-to-SLB or SLB-to-FWLB. The configuration commands after the figure show how to configure SLB-to-FWLB.

Figure 8.3 Active-Active FWLB with SLB

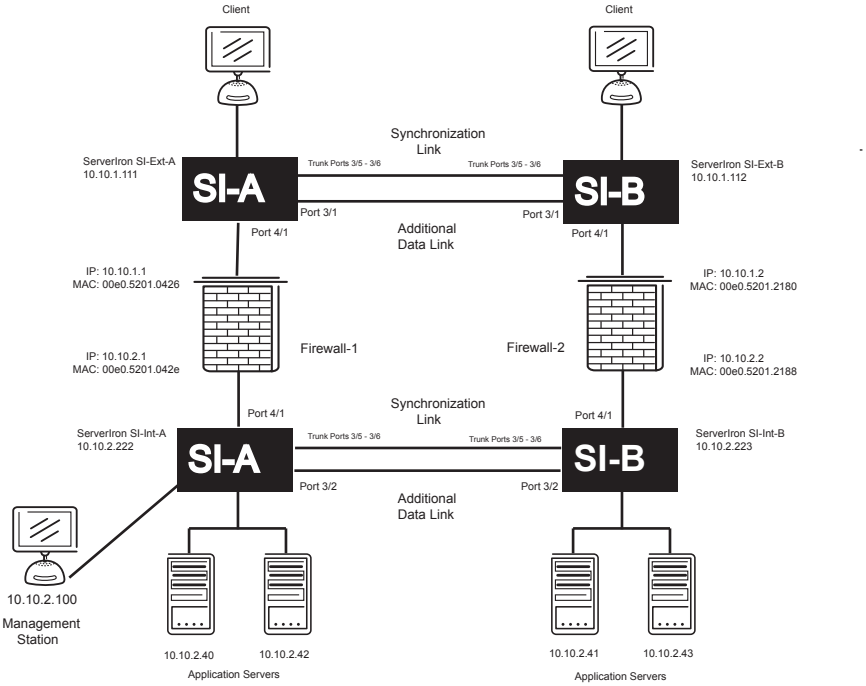

The CLI commands in this section show how to configure SLB-to-FWLB. In SLB-to-FWLB, the ServerIron on the Internet side of the firewalls performs FWLB for traffic directed toward real servers connected to the ServerIron on the private side of the firewalls. The real servers are configured as remote servers. In addition, the SLB-to-FWLB feature is enabled on the Internet ServerIron. The internal ServerIron is configured for FWLB but requires no additional configuration.

Commands on External ServerIron A (SI-Ext-A)

The following commands change the CLI to the global CONFIG level, then change the hostname to "SI-Ext-A".

```
ServerIron> enable
ServerIron# configure terminal
ServerIron(config)# hostname SI-Ext-A
```

The following command enable the always-active feature and disables the Spanning Tree Protocol (STP) in VLAN 1, which contains the ports that will carry the FWLB traffic.

```
SI-Ext-A(config)# vlan 1
SI-Ext-A(config-vlan-1)# always-active
SI-Ext-A(config-vlan-1)# no spanning-tree
```

The following commands configure a virtual routing interface on VLAN 1 (the default VLAN), then configure an IP address on the interface. The virtual routing interface is associated with all the ports in the VLAN.

```
SI-Ext-A(config-vlan-1)# router-interface ve 1
SI-Ext-A(config-vlan-1)# exit
SI-Ext-A(config)# interface ve 1
SI-Ext-A(config-ve-1)# ip address 10.10.1.111 255.255.255.0
SI-Ext-A(config-ve-1)# exit
```

The following command configures an IP default route. The next hop for this route is the ServerIron's interface with firewall FW1.

```
SI-Ext-A(config)# ip route 0.0.0.0 0.0.0.0 10.10.1.1
```

The following commands configure the dedicated synchronization link between the ServerIron and its active-active partner. The **trunk** command configures the two ports of the link into a trunk group. The next two commands add the trunk group to a separate port-based VLAN, since the synchronization link must be in its own VLAN. The

server fw-port command identifies the port number the link is on. If the link is a trunk group, you must specify the MAC address of the group's primary port.

```
SI-Ext-A(config)# trunk switch ethernet 3/5 to 3/6
SI-Ext-A(config)# vlan 10
SI-Ext-A(config-vlan-10)# untagged ethernet 3/5 to 3/6
SI-Ext-A(config-vlan-10)# exit
SI-Ext-A(config)# server fw-port 3/5
```

The following command configures the data link between this ServerIron and its active-active partner. You must use the **server partner-ports** command to specify all the data links with the partner. However, do not use the command for the synchronization link.

NOTE: The **server partner-ports** command is required for all IronClad FWLB configurations in software release 08x.

```
SI-Ext-A(config)# server partner-ports ethernet 3/1
```

The following commands add the firewall definitions. In this example, port HTTP is specified for each firewall. Specifying the application ports on the firewalls is optional. The **port http no-health-check** command under each firewall disables the Layer 4 health check for the HTTP port. When you add an application port to a firewall definition, the ServerIron automatically enables the Layer 4 health check for that port. You must disable the Layer 4 health check if the firewall is unable to act as a proxy for the application and respond to the health check. If the firewall does not respond to the health check, the ServerIron assumes that the port is unavailable and stops sending traffic for the port to the firewall.

The ServerIron will still use a Layer 3 health check (IP ping) to test connectivity to the firewall.

```
SI-Ext-A(config)# server fw-name fw1 10.10.1.1
SI-Ext-A(config-rs-fw1)# port http
SI-Ext-A(config-rs-fw1)# port http no-health-check
SI-Ext-A(config-rs-fw1)# exit
SI-Ext-A(config)# server fw-name fw2 10.10.1.2
SI-Ext-A(config-rs-fw2)# port http
SI-Ext-A(config-rs-fw2)# port http no-health-check
SI-Ext-A(config-rs-fw2)# exit
```

The following commands add the firewall definitions to the firewall port group (always group 2). The firewall group contains all the ports in VLAN 1 (the default VLAN).

```
SI-Ext-A(config)# server fw-group 2
SI-Ext-A(config-tc-2)# fw-name fw1
SI-Ext-A(config-tc-2)# fw-name fw2
```

The following command enables the active-active mode.

```
SI-Ext-A(config-tc-2)# sym-priority 1
```

NOTE: Do not use the same number on both ServerIrons. For example, use enter **sym-priority 1** on one of the ServerIrons and **sym-priority 255** on the other ServerIron.

The following commands add the paths through the firewalls to the other ServerIron. Each path consists of a path number, a ServerIron port number, the IP address at the other end of the path, and the next-hop IP address. In this example, the topology does not contain routers other than the ServerIrons. If your topology does contain other routers, configure firewall paths for the routers too. For router paths, use the same IP address as the path destination and the next hop.

NOTE: The path IDs must be in contiguous, ascending numerical order, starting with 1. For example, path sequence 1, 2, 3 , 4 is valid. Path sequence 4, 3, 2, 1 or 1, 3, 4, 5 is not valid.

```
SI-Ext-A(config-tc-2)# fwall-info 1 4/1 10.10.2.222 10.10.1.1
SI-Ext-A(config-tc-2)# fwall-info 2 3/1 10.10.2.222 10.10.1.2
SI-Ext-A(config-tc-2)# fwall-info 3 4/1 10.10.2.223 10.10.1.1
SI-Ext-A(config-tc-2)# fwall-info 4 3/1 10.10.2.223 10.10.1.2
```


The following command sets the load balancing method to balance requests based on the firewall that has the least number of connections for the requested service. Since the firewall definitions above specify the HTTP service, the ServerIron will load balance requests based on the firewall that has fewer HTTP session entries in the ServerIron session table.

```
SI-Ext-A(config-tc-2)# fw-predictor per-service-least-conn
```

The following command is part of the always-active feature, which provides the additional data link between the this ServerIron and its partner.

```
SI-Ext-A(config-tc-2)# l2-fwall
SI-Ext-A(config-tc-2)# exit
```

The following commands add static MAC entries for the firewall interfaces with the ServerIron. The static MAC entries are required only if the configuration uses static routes and a single virtual routing interface, as in this example, and if the default gateway for the client or server is the firewall. If the configuration uses a dynamic routing protocol (for example, RIP or OSPF), the static entries are not required. Alternatively, the static entries are not required if you use the ServerIron itself as the default gateway for the client or the server. For example, the static entries are not required if you configure the client to use 10.10.1.111 as its default gateway.

```
SI-Ext-A(config)# vlan 1
SI-Ext-A(config-vlan-1)# static-mac-address 00e0.5201.0426 ethernet 4/1
priority 1 router-type
SI-Ext-A(config-vlan-1)# static-mac-address 00e0.5203.2f80 ethernet 3/1
priority 1 router-type
SI-Ext-A(config-vlan-1)# exit
```

The following commands configure the SLB parameters, four real servers and one VIP. The servers are bound to the VIP by the HTTP port. Notice that the servers are configured as remote servers. If Proxy ARP is enabled on the internal ServerIrons, you can define the real servers as local servers instead of remote servers. However, if Proxy ARP is not enabled on the internal ServerIrons, the real servers must be remote servers.

```
SI-Ext-A(config)# server remote-name web1 10.10.2.40
SI-Ext-A(config-rs-web1)# port http
SI-Ext-A(config-rs-web1)# server remote-name web2 10.10.2.41
SI-Ext-A(config-rs-web2)# port http
SI-Ext-A(config-rs-web2)# server remote-name web3 10.10.2.42
SI-Ext-A(config-rs-web3)# port http
SI-Ext-A(config-rs-web3)# server remote-name web4 10.10.2.43
SI-Ext-A(config-rs-web4)# port http
SI-Ext-A(config-rs-web4)# server virtual webby 10.10.1.10
SI-Ext-A(config-vs-webby)# port http
SI-Ext-A(config-vs-webby)# bind http web4 http web3 http web2 http web1 http
```

Enter the following command to enable SLB-to-FWLB.

NOTE: This command applies only to the ServerIrons that contain the SLB configuration. Do not enter this command on the internal ServerIrons.

```
SI-Ext-A(config)# server slb-fw
```

The following commands assign FWLB processing for all forwarding modules to the same WSM CPU. The device uses the same CPU to process all FWLB traffic. You must assign all the traffic to the same WSM CPU. The commands in this example assign traffic on the forwarding modules in slots 3 and 4 to WSM CPU 1 on the Web Switching Management Module in slot 2.

```
SI-Ext-A(config)# wsm wsm-map slot 3 wsm-slot 2 wsm-cpu 1
SI-Ext-A(config)# wsm wsm-map slot 4 wsm-slot 2 wsm-cpu 1
```

NOTE: For simplicity, the configuration of the other ServerIrons in this example do not include **wsm wsm-map** commands. However, the commands you need to enter depend on the slot locations of the modules in the device and the WSM CPU you want to use.

The following commands enable FWLB.


```
SI-Ext-A(config)# ip l4-policy 1 fw tcp 0 global
SI-Ext-A(config)# ip l4-policy 2 fw udp 0 global
```

The following command saves the configuration changes to the startup-config file.

```
SI-Ext-A(config)# write memory
```

Commands on External ServerIron B (SI-Ext-B)

Here are the commands for configuring SI-Ext-B in Figure 8.3 on page 8-12. The SLB configuration is identical to the one on SI-Ext-A.

```
ServerIron> enable
ServerIron# configure terminal
ServerIron(config)# hostname SI-Ext-B
SI-Ext-B(config)# vlan 1
SI-Ext-B(config-vlan-1)# always-active
SI-Ext-B(config-vlan-1)# no spanning-tree
SI-Ext-B(config-vlan-1)# router-interface ve 1
SI-Ext-B(config-vlan-1)# exit
SI-Ext-B(config)# interface ve 1
SI-Ext-B(config-ve-1)# ip address 10.10.1.112 255.255.255.0
SI-Ext-B(config-ve-1)# exit
SI-Ext-B(config)# ip route 0.0.0.0 0.0.0.0 10.10.1.1
SI-Ext-B(config)# trunk switch ethernet 3/5 to 3/6
SI-Ext-B(config)# vlan 10
SI-Ext-B(config-vlan-10)# untagged ethernet 3/5 to 3/6
SI-Ext-B(config-vlan-10)# exit
SI-Ext-B(config)# server fw-port 3/5
SI-Ext-B(config)# server partner-ports ethernet 3/1
SI-Ext-B(config)# server fw-name fw1 10.10.1.1
SI-Ext-B(config-rs-fw1)# port http
SI-Ext-B(config-rs-fw1)# port http no-health-check
SI-Ext-B(config-rs-fw1)# exit
SI-Ext-B(config)# server fw-name fw2 10.10.1.2
SI-Ext-B(config-rs-fw2)# port http
SI-Ext-B(config-rs-fw2)# port http no-health-check
SI-Ext-B(config-rs-fw2)# exit
SI-Ext-B(config)# server fw-group 2
SI-Ext-B(config-tc-2)# fw-name fw1
SI-Ext-B(config-tc-2)# fw-name fw2
SI-Ext-B(config-tc-2)# sym-priority 255
SI-Ext-B(config-tc-2)# fwall-info 1 3/1 10.10.2.222 10.10.1.1
SI-Ext-B(config-tc-2)# fwall-info 2 4/1 10.10.2.222 10.10.1.2
SI-Ext-B(config-tc-2)# fwall-info 3 3/1 10.10.2.223 10.10.1.1
SI-Ext-B(config-tc-2)# fwall-info 4 4/1 10.10.2.223 10.10.1.2
SI-Ext-B(config-tc-2)# fw-predictor per-service-least-conn
SI-Ext-B(config-tc-2)# l2-fwall
SI-Ext-B(config-tc-2)# exit
SI-Ext-B(config)# vlan 1
SI-Ext-B(config-vlan-1)# static-mac-address 00e0.5201.0426 ethernet 3/1
priority 1 router-type
SI-Ext-B(config-vlan-1)# static-mac-address 00e0.5203.2f80 ethernet 4/1
priority 1 router-type
SI-Ext-B(config-vlan-1)# exit
SI-Ext-B(config)# server remote-name web1 10.10.2.40
SI-Ext-B(config-rs-web1)# port http
SI-Ext-B(config-rs-web1)# server remote-name web2 10.10.2.41
SI-Ext-B(config-rs-web2)# port http
SI-Ext-B(config-rs-web2)# server remote-name web3 10.10.2.42
SI-Ext-B(config-rs-web3)# port http
```



```
SI-Ext-B(config-rs-web3)# server remote-name web4 10.10.2.43
SI-Ext-B(config-rs-web4)# port http
SI-Ext-B(config-rs-web4)# server virtual webby 10.10.1.10
SI-Ext-B(config-vs-webby)# port http
SI-Ext-B(config-vs-webby)# bind http web4 http web3 http web2 http web1 http
SI-Ext-B(config)# server slb-fw
SI-Ext-B(config)# ip l4-policy 1 fw tcp 0 global
SI-Ext-B(config)# ip l4-policy 2 fw udp 0 global
SI-Ext-B(config)# write memory
```

Commands on Internal ServerIron A (SI-Int-A)

```
ServerIron> enable
ServerIron# configure terminal
ServerIron(config)# hostname SI-Int-A
SI-Int-A(config)# vlan 1
SI-Int-A(config-vlan-1)# always-active
SI-Int-A(config-vlan-1)# no spanning-tree
SI-Int-A(config-vlan-1)# router-interface ve 1
SI-Int-A(config-vlan-1)# exit
SI-Int-A(config)# interface ve 1
SI-Int-A(config-ve-1)# ip address 10.10.2.222 255.255.255.0
SI-Int-A(config-ve-1)# exit
SI-Int-A(config)# ip route 0.0.0.0 0.0.0.0 10.10.2.1
SI-Int-A(config)# trunk switch ethernet 3/5 to 3/6
SI-Int-A(config)# vlan 10
SI-Int-A(config-vlan-10)# untagged ethernet 3/5 to 3/6
SI-Int-A(config-vlan-10)# exit
SI-Int-A(config)# server fw-port 3/5
SI-Int-A(config)# server partner-ports ethernet 3/2
SI-Int-A(config)# server fw-name fw1 10.10.2.1
SI-Int-A(config-rs-fw1)# port http
SI-Int-A(config-rs-fw1)# port http no-health-check
SI-Int-A(config-rs-fw1)# exit
SI-Int-A(config)# server fw-name fw2 10.10.2.2
SI-Int-A(config-rs-fw2)# port http
SI-Int-A(config-rs-fw2)# port http no-health-check
SI-Int-A(config-rs-fw2)# exit
SI-Int-A(config)# server fw-group 2
SI-Int-A(config-tc-2)# fw-name fw1
SI-Int-A(config-tc-2)# fw-name fw2
SI-Int-A(config-tc-2)# sym-priority 1
SI-Int-A(config-tc-2)# fwall-info 1 4/1 10.10.1.111 10.10.2.1
SI-Int-A(config-tc-2)# fwall-info 2 3/2 10.10.1.111 10.10.2.2
SI-Int-A(config-tc-2)# fwall-info 3 4/1 10.10.1.112 10.10.2.1
SI-Int-A(config-tc-2)# fwall-info 4 3/2 10.10.1.112 10.10.2.2
SI-Int-A(config-tc-2)# fw-predictor per-service-least-conn
SI-Int-A(config-tc-2)# l2-fwall
SI-Int-A(config-tc-2)# exit
SI-Int-A(config)# vlan 1
SI-Int-A(config-vlan-1)# static-mac-address 00e0.5201.042e ethernet 4/1
priority 1 router-type
SI-Int-A(config-vlan-1)# static-mac-address 00e0.5201.2f88 ethernet 3/2
priority 1 router-type
SI-Int-A(config-vlan-1)# exit
SI-Int-A(config)# ip l4-policy 1 fw tcp 0 global
SI-Int-A(config)# ip l4-policy 2 fw udp 0 global
SI-Int-A(config)# write memory
```


Commands on Internal ServerIron B (SI-Int-B)

```
ServerIron> enable
ServerIron# configure terminal
ServerIron(config)# hostname SI-Int-B
SI-Int-B(config)# vlan 1
SI-Int-B(config-vlan-1)# always-active
SI-Int-B(config-vlan-1)# no spanning-tree
SI-Int-B(config-vlan-1)# router-interface ve 1
SI-Int-B(config-vlan-1)# exit
SI-Int-B(config)# interface ve 1
SI-Int-B(config-ve-1)# ip address 10.10.2.223 255.255.255.0
SI-Int-B(config-ve-1)# exit
SI-Int-B(config)# ip route 0.0.0.0 0.0.0.0 10.10.2.2
SI-Int-B(config)# trunk switch ethernet 3/5 to 3/6
SI-Int-B(config)# vlan 10
SI-Int-B(config-vlan-10)# untagged ethernet 3/5 to 3/6
SI-Int-B(config-vlan-10)# exit
SI-Int-B(config)# server fw-port 3/5
SI-Int-B(config)# server partner-ports ethernet 3/2
SI-Int-B(config)# server fw-name fw1 10.10.2.1
SI-Int-B(config-rs-fw1)# port http
SI-Int-B(config-rs-fw1)# port http no-health-check
SI-Int-B(config-rs-fw1)# exit
SI-Int-B(config)# server fw-name fw2 10.10.2.2
SI-Int-B(config-rs-fw2)# port http
SI-Int-B(config-rs-fw2)# port http no-health-check
SI-Int-B(config-rs-fw2)# exit
SI-Int-B(config)# server fw-group 2
SI-Int-B(config-tc-2)# fw-name fw1
SI-Int-B(config-tc-2)# fw-name fw2
SI-Int-B(config-tc-2)# sym-priority 255
SI-Int-B(config-tc-2)# fwall-info 1 3/2 10.10.1.111 10.10.2.1
SI-Int-B(config-tc-2)# fwall-info 2 4/1 10.10.1.111 10.10.2.2
SI-Int-B(config-tc-2)# fwall-info 3 3/2 10.10.1.112 10.10.2.1
SI-Int-B(config-tc-2)# fwall-info 4 4/1 10.10.1.112 10.10.2.2
SI-Int-B(config-tc-2)# fw-predictor per-service-least-conn
SI-Int-B(config-tc-2)# l2-fwall
SI-Int-B(config-tc-2)# exit
SI-Int-B(config)# vlan 1
SI-Int-B(config-vlan-1)# static-mac-address 00e0.5201.042e ethernet 3/2
priority 1 router-type
SI-Int-B(config-vlan-1)# static-mac-address 00e0.5201.2f88 ethernet 4/1
priority 1 router-type
SI-Int-B(config-vlan-1)# exit
SI-Int-B(config)# ip l4-policy 1 fw tcp 0 global
SI-Int-B(config)# ip l4-policy 2 fw udp 0 global
SI-Int-B(config)# write memory
```

Chapter 9
Viewing FWLB Configuration Details and Statistics

You can view the following FWLB configuration details and statistics:

- Firewall group information – Displays the firewall configuration, the status of each firewall, and traffic statistics for traffic between each firewall and the ServerIron.
- Firewall path information – Shows the synchronization paths configured for the firewall.

  NOTE: The information is shown from this ServerIron's perspective. To view the other side of the path configuration, display the firewall path information on the ServerIron at the other end of the path.

- Hashing information – Shows the firewall selected by the hashing algorithm for a given pair of source and destination addresses.

Displaying Firewall Group Information

To display configuration information, state information, and traffic statistics for the firewall group, use the following CLI method.

USING THE CLI

To access FWLB configuration details and statistics, enter the following command at any level of the CLI:

```
ServerIron(config)# show fw-group
```

This command shows the following information. To explain the output, this example is divided into sections for discussion. The output is not divided in this way on the screen of your management terminal.

The first line indicates the firewall group and the number of firewalls in the group. This firewall group is group number 2 and contains two firewalls. The second line shows the source and destination values for the hash mask.

```
 Firewall-group 2 has 2 members Admin-status = Enabled
 Hash_info: Dest_mask = 255.255.255.255 Src_mask = 255.255.255.255
```

The following lines list the firewalls configured in the firewall group, show the administrative state, and have distribution values for each firewall.

- The administrative state is shown in the Admin-st column and depends on the results of the Layer 3 health check (ping) the ServerIron performs when you add the path information for the firewall. The administrative state can be one of the following:
  - 0 – Disabled
  - 1 – Enabled

- 2 – Failed
- 3 – Testing
- 4 – Suspect
- 6 – Active

NOTE: Status 5 (Graceful Shutdown) does not apply to firewalls.

- The Hash-distribution field shows how many hash values are assigned to the server. This information is relevant only when no path information is configured for the firewall group. If the group is using paths, the hash-distribution value is always 0.

```
Firewall Server Name             Admin-st  Hash-distribution
        fw1                                      1                  0
        fw2                                      6                  0
```

The following lines show traffic statistics for each firewall. The Name field lists the name of the firewall and the IP field shows the IP address of the firewall. The "Host" indicates the ServerIron. The “Firewall” indicates the firewall. The Groups field shows the firewall group number.

The statistics are for traffic between the ServerIron and the firewall. The CurConn and TotConn columns show the total number of connections between the ServerIron and the firewall. A connection represents both send and receive traffic. (Thus, each connection shown here is equivalent to two sessions.) The Packets and Octets fields show the total number of packets and octets exchanged by the ServerIron and the firewall.

```
Traffic From<->to Firewall Servers
====================================

Name: fw1             IP: 10.10.0.1        State: 1   Groups =   2

                                          Host->Firewall         Firewall->Host
              State CurConn TotConn    Packets     Octets    Packets     Octets
Firewall     active       0       0          0          0          0          0
Total                     0       0          0          0          0          0

Name: fw2             IP: 10.10.0.2        State: 6   Groups =   2

                                          Host->Firewall         Firewall->Host
              State CurConn TotConn    Packets     Octets    Packets     Octets
Firewall     active       0       0          0          0          0          0
Total                     0       0          0          0          0          0
```

TCP/UDP Port Statistics

If you associated TCP/UDP application ports with specific firewalls (part of a stateful FWLB configuration), rows of statistics for the application ports also are listed. The following example shows statistics for two ServerIrons in a

basic stateful FWLB configuration. In this example, HTTP traffic and Telnet traffic are explicitly associated with fw1 and fw2.

```
ServerIronA(config)# show fw-group
Firewall-group 2 has 2 members Admin-status = Enabled Active = 0
Hash_info: Dest_mask = 255.255.255.255 Src_mask = 255.255.255.255

Firewall Server Name             Admin-status Hash-distribution
fw1-IPin                         6            0
fw2-IPin                         6            0

Traffic From<->to Firewall Servers

Name: fw1-IPin         IP: 209.157.22.3     State: 6   Groups =   2

                                    Host->Fw-Server       Fw-Server->Host
           State   CurConn TotConn Packets    Octets     Packets    Octets
http       active  8       37445   264015     18232357   326241     339770826
In-http    active  0       0       0          0          0          0
telnet     active  0       0       0          0          0          0
In-telnet  active  0       0       0          0          0          0
Fw-Server  active  0       0       0          0          21         2384
Total              8       37445   264015     18232357   326262     339773210

Name: fw2-IPin       IP: 209.157.22.4       State: 6   Groups =   2

                                    Host->Fw-Server       Fw-Server->Host
           State   CurConn TotConn Packets    Octets     Packets    Octets
http       active  8       30655   216957     14977685   110028     114839282
In-http    active  0       0       0          0          0          0
telnet     active  0       0       0          0          0          0
In-telnet  active  0       0       0          0          0          0
Fw-Server  active  0       0       0          0          25         2912
Total              8       30655   216957     14977685   110053     114842194
```

The statistics for TCP/UDP traffic to and from the ServerIron are highlighted in bold type in this example. The “http” and “telnet” rows show statistics for traffic initiated by clients or servers. The “In-http” and “In-telnet” rows show statistics for replies. In the example shown above, the statistics indicate requests from clients outside the firewalls sent to servers on the private side of the firewalls.

In general, a ServerIron will show statistics for only one direction:

- If the ServerIron is on the external (Internet) side of the firewalls, the ServerIron will show statistics in the “http” row, “telnet” row, and so on. For example, statistics for STP SYN packets from clients are listed in the “http” row.
- If the ServerIron is on the internal (private network) side of the firewalls, the ServerIron will show statistics in the “In-http” row, “In-telnet” row, and so on. For example, server replies to TCP SYN packets from clients are listed in the “In-http” row.

The example above is for the external ServerIron (ServerIron A). The following example shows statistics for the internal ServerIron (ServerIron B).

```
ServerIronB(config)# show fw-group
Firewall-group 2 has 2 members Admin-status = Enabled Active = 0
Hash_info: Dest_mask = 255.255.255.255 Src_mask = 255.255.255.255

Firewall Server Name             Admin-status Hash-distribution
fw1-IPout                        6            0
fw2-IPout                        6            0

Traffic From<->to Firewall Servers
Name: fw1-IPout        IP: 209.157.23.1     State: 6   Groups =   2
                                 Host->Fw-Server       Fw-Server->Host
           State   CurConn TotConn Packets    Octets     Packets    Octets
http       active  0       0       0          0          0          0
In-http    active  3       11118   71054      74037240   78564      5422929
Fw-Server  active  0       0       0          0          0          0
Total              3       11118   71054      74037240   78564      5422929

Name: fw2-IPout        IP: 209.157.23.2     State: 6   Groups =   2
                                 Host->Fw-Server       Fw-Server->Host
           State   CurConn TotConn Packets    Octets     Packets    Octets
http       active  0       0       0          0          0          0
In-http    active  4       9182    59169      61874490   65057      4490977
Fw-Server  active  0       0       0          0          0          0
Total              4       9182    59169      61874490   65057      4490977
```

In this example, the ServerIron shows statistics for server replies to client requests. The **show fw-group** command on the external ServerIron (ServerIron A) shows the requests, while the **show fw-group** command on the internal ServerIron (ServerIron B), shows the server replies to those requests.

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write or read-only access.
2. Click on the plus sign next to Monitor in the tree view to expand the list of monitoring options.
3. Click on the plus sign next to FWLB in the tree view to expand the list of Firewall Load Balancing option links.
4. Select the Firewall Traffic link. See the section above for descriptions of the information shown by this display.

Displaying Firewall Path Information

The ServerIron uses paths that you configure to provide synchronization for the traffic that passes through the ServerIrons and the Layer 3 firewalls between them. You can display configuration information, state information, and statistics for the paths using the following CLI method.

USING THE CLI

To display path information for FWLB, enter the following command at any level of the CLI:

```
ServerIron(config)# show server fw-path
 Firewall Server Path Info
 Number of Fwall = 4

Target-ip      Next-hop-ip    Port Path Status  Tx Rx State
3.3.3.10       1.1.1.3        1    1    1       1  1  3
3.3.3.10       1.1.1.4        2    2    1       1  1  3
4.4.4.10       1.1.1.3        1    3    1       1  1  5
4.4.4.10       1.1.1.4        2    4    1       1  1  5

 State = 5 :Partner known = No :port = 1
 Activates = 2 Inactivations = 2

                  Current     Local   Partner
State                   5         5         5
Priority                5         5        10
Path-cnt                4         4         4

 Active path cnt = 2, list =   3   4
```

The following table describes the information displayed by the **show server fw-path** command.

Table 9.1: FWLB Path Information

This Field...	Displays...
General Information	
Number of Fwall	The number of firewalls configured in the group.
Target-ip	The IP address of the device at the other end of the path. Generally, this other device is another ServerIron.
Next-hop-ip	The IP address of the device at the next hop to the target IP. Usually, this is the IP interface on the firewall that is connected to this ServerIron.
Port	The ServerIron port for this path. This is the port connected to the firewall.
Path	The path ID.
Status	The status of the path, which can be one of the following: • 0 – The path is down. • 1 – The path is up.
Tx	Indicates the state of the transmit side of the path. The state can be one of the following: • 0 – The transmit side is down. • 1 – The transmit side is up.

Table 9.1: FWLB Path Information (Continued)

This Field...	Displays...
Rx	Indicates the state of the receive side of the path. The state can be one of the following: • 0 – The receive side is down. • 1 – The receive side is up.
State	The state of the other end of the path. The state can be one of the following: • 3 – The ServerIron at the other end of the path is in standby mode for the firewall group. • 5 – The ServerIron at the other end of the path is in active mode for the firewall group. **Note**: This field applies only to IronClad FWLB. If the ServerIron is not configured with another ServerIron as the active or backup ServerIron for IronClad FWLB, the state is always 0.
IronClad FWLB Information **Note**: These fields apply only to IronClad FWLB.	
State	The state of this end of the path. The state can be one of the following: • 0 – Unknown. Generally, this indicates that the link is down. • 3 – The ServerIron is in standby mode for the firewall group. • 5 – The ServerIron is in active mode for the firewall group. **Note**: This field applies only to IronClad FWLB. If the ServerIron is not configured with another ServerIron as the active or backup ServerIron for IronClad FWLB, the state is always 0.
Partner known	Indicates whether this ServerIron can see (has Layer 2 connectivity with) the other ServerIron in the pair. This field can have one of the following values: • No – This ServerIron does not have Layer 2 connectivity with its partner. Generally, this indicates that the link is down. • Yes – This ServerIron has Layer 2 connectivity with its partner. **Note**: This field applies only to the other ServerIron in an active-standby configuration for IronClad FWLB.
Port	The port that connects this ServerIron to its partner. If the partners are connected by a trunk group, this port number is the number of the primary port (the lowest numbered port) in the trunk group.
Activates	The number of times this ServerIron has changed from being the standby ServerIron to become the active ServerIron for the firewall group.
Inactivations	The number of times this ServerIron has changed from being the active ServerIron to become the standby ServerIron for the firewall group.

Table 9.1: FWLB Path Information (Continued)

This Field...	Displays...
State information for IronClad FWLB The Current, Local, and Partner columns show the following: • Current shows the immediate state information. • Local shows the normalized state information. When the current state remains unchanged for three seconds, the current state value becomes the local state value. Local state information is used to compute the active-standby status. The local state is usually the same as the current state; however, the local state is a normalized version of the current state. • Partner shows the state on the other ServerIron in the pair. **Note**: These fields apply only to IronClad FWLB.	
State	Current, local, and active state information for the path. • The current state indicates the immediate state information. This is the most current information. • The local state indicates the cumulative current states over a three-second interval. If the current states have been the same for the previous three seconds, the state is shown in the Local column. • The partner state. In each column, the state can be one of the following: • 0 – Unknown. Generally, this indicates that the link is down. • 3 – The ServerIron is in standby mode for the firewall group. • 5 – The ServerIron is in active mode for the firewall group.
Priority	The IronClad FWLB priority for the firewalls in the firewall group. The ServerIron with the higher priority for the group ID the default active ServerIron for the group.
Path-cnt	The number of firewall paths.
Active path cnt	The number of paths from this ServerIron that go to active ServerIrons. A path that goes to a ServerIron that is in standby mode is not counted in this statistic.
list	The path IDs of the active paths.

USING THE WEB MANAGEMENT INTERFACE

1. Log on to the device using a valid user name and password for read-write or read-only access.
2. Click on the plus sign next to Monitor in the tree view to expand the list of monitoring options.
3. Click on the plus sign next to FWLB in the tree view to expand the list of Firewall Load Balancing option links.
4. Select the Firewall Path link. See the section above for descriptions of the information shown by this display.

Displaying the Firewall Selected by the Hashing

Process for Load Balancing

By default, FWLB uses a hashing algorithm to select a firewall for a packet based on the packet's source and destination IP address. Optionally, you can configure the ServerIron to also hash based on source and destination TCP or UDP application ports. Once the ServerIron selects a firewall for a given pair of source and destination IP addresses (and, if specified, source and destination TCP or UDP application ports), the ServerIron always selects the same firewall for packets with the same address pairs.

To display the firewall that the hashing algorithm selected for a given pair of source and destination addresses, enter a command such as the following:

```
ServerIron# show fw-hash 1.1.1.1 2.2.2.2 2
fw3
```

In this example, the command output indicates that the FWLB hashing algorithm selected firewall "fw3" for traffic to IP address 1.1.1.1 from IP address 2.2.2.2.

Syntax: show fw-hash <dst-ip-addr> <src-ip-addr> <fwall-group-id>
[<protocol> <dst-tcp/udp-port> <src-tcp/udp-port>]

The <dst-ip-addr> parameter specifies the destination IP address.

The <src-ip-addr> parameter specifies the source IP address.

The <fwall-group-id> parameter specifies the FWLB group ID. Normally, the FWLB group ID is 2.

The <protocol> parameter specifies the protocol number for TCP or UDP. You can specify one of the following:

- 6 – TCP
- 17 – UDP

The <dst-tcp/udp-port> specifies the destination TCP or UDP application port number.

The <src-tcp/udp-port> specifies the source TCP or UDP application port number.

If you configured the ServerIron to hash based on source and destination TCP or UDP application ports as well as IP addresses, the ServerIron might select more than one firewall for the same pair of source and destination IP addresses, when the traffic uses different pairs of source and destination application ports. Use the optional parameters to ensure that the command's output distinguishes among the selected firewalls based on the application ports. Here is an example:

```
ServerIron# show fw-hash 1.1.1.1 2.2.2.2 2 6 80 8080
fw2
ServerIron# show fw-hash 1.1.1.1 2.2.2.2 2 6 80 9000
fw3
```


Chapter 10
Configuring FWLB for Layer 2 Firewalls

The steps for configuring IronClad FWLB for Layer 2 firewalls are similar to those for configuring Layer 3 FWLB for static routes. In addition to the basic FWLB configurations steps, perform the following steps:

- On each ServerIron, configure all the ports connected to all the firewalls as a trunk group.
- Disable the Spanning Tree Protocol (STP). STP is enabled by default on the ServerIron.
- Disable Layer 2 traffic on the standby ServerIrons. To do so, you specify the L2-fwall option on each ServerIron (both active and standby). This step is required because all traffic on the standby firewalls in a static route configuration must be blocked. Normally, the standby ServerIron blocks only routing protocol packets but allows other types of packets to pass through the device. In a static route configuration, you need to block all the traffic from passing through the standby ServerIron. If you do not enable this mode to block the traffic, loops can occur.

  You must enable the L2-fwall option on all the ServerIron in the configuration, whether they are active or standby by default.

In addition, when you configure the paths through the firewalls to the other ServerIrons, you do not specify the firewall IP address as the next hop. Instead, you specify the IP address of the other ServerIron as the path's next hop, as well as the path destination.

Configuring FWLB for Layer 2 Firewalls

Figure 10.1 on page 10-2 shows an example of an IronClad FWLB configuration for Layer 2 firewalls.

NOTE: This example is for an IronClad configuration. However, you also can configure ServerIrons for basic FWLB with Layer 2 firewalls.

As shown in this example, the Internet router has two static routes. One of the static routes goes to the router interface connected to the internal ServerIrons on the other side of the firewalls. The other static route goes to the network on the other side of the internal network router.

The internal network router has a default route that goes to the IP interface on the Internet router that is connected to the ServerIrons.

The IP interface on each router that is connected to ServerIrons is a virtual interface. On Brocade Layer 3 Switchs, you can configure the same IP address on multiple ports if you configure the ports in a port-based VLAN, add a virtual interface to the VLAN, and then configure the IP address on the virtual interface. The configuration example at the end of this section includes the CLI commands for configuring the interface.

The default gateway for each ServerIron is its local router interface.

Figure 10.1 IronClad Layer 2 FWLB configuration

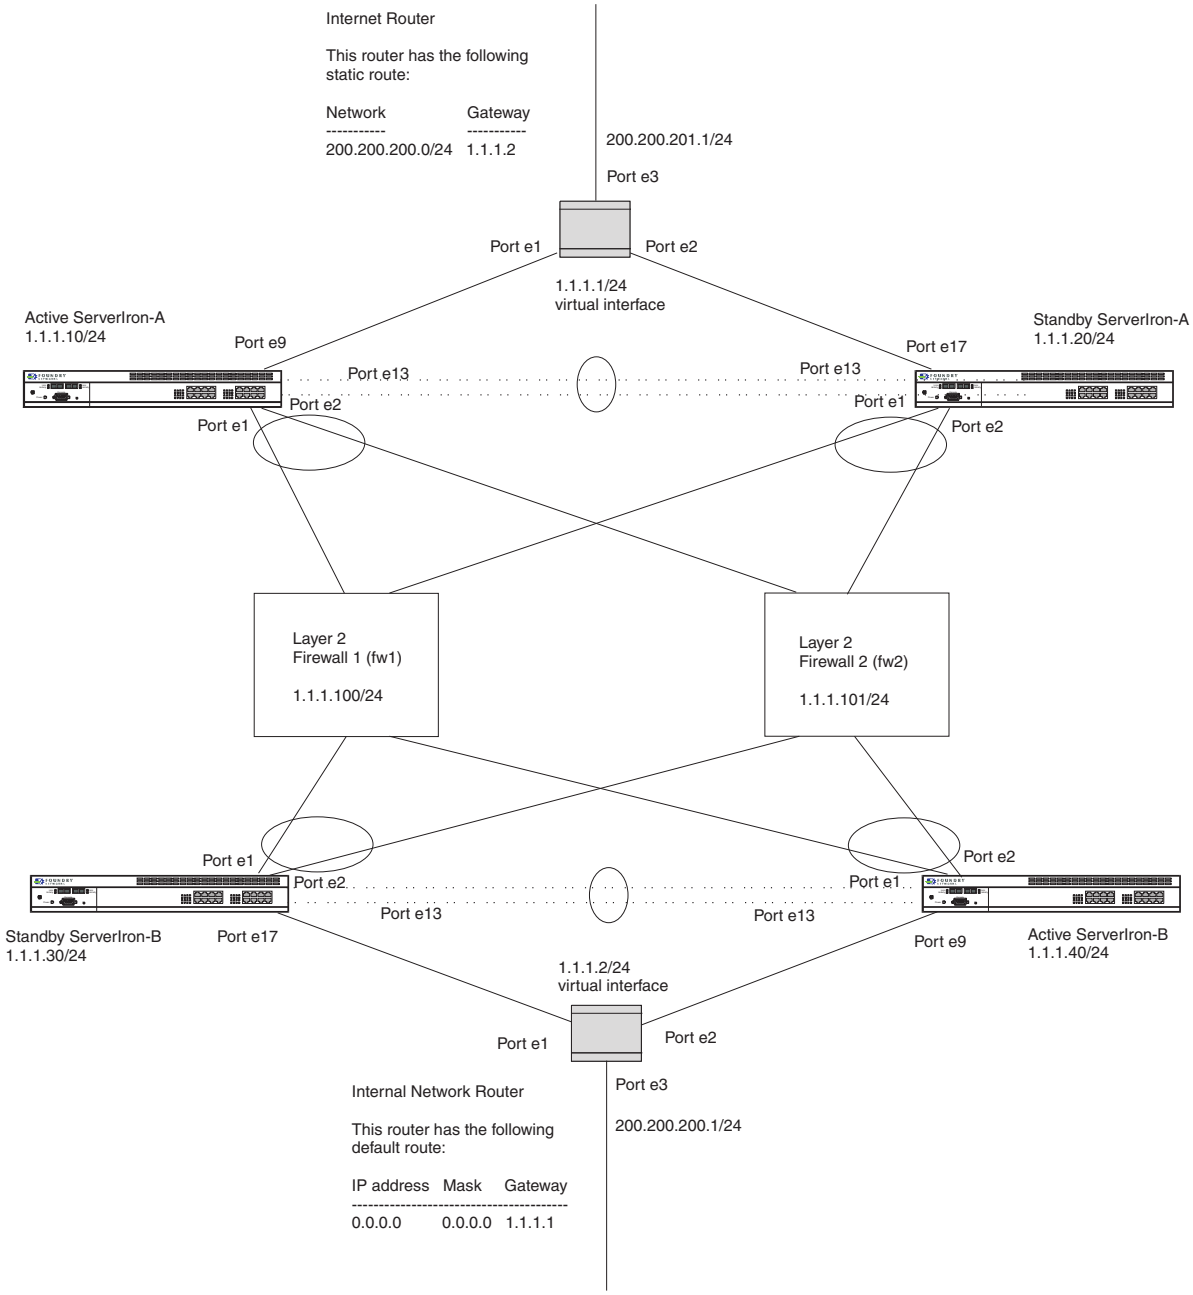

To configure IronClad FWLB for Layer 2 firewalls, perform the following tasks.

Table 10.1: Configuration tasks – IronClad FWLB for Layer 2 firewalls

Task	See page...
Configure Global Parameters	
Configure a switch trunk group for all the ports connected to the firewalls.	10-3

Table 10.1: Configuration tasks – IronClad FWLB for Layer 2 firewalls(Continued)

Task	See page...
Identify the partner port (the link between the active and standby ServerIrons)	10-3
Identify the router port (ServerIron ports connected to routers)	10-4
Configure Firewall Parameters	
Define the firewalls and add them to the firewall group	10-4
Configure Firewall Group Parameters	
Enable the L2-fwall option	10-5
Configure the paths and add static MAC entries for the firewall interfaces with the ServerIron	10-5
Specify the ServerIron priority (determines which ServerIron in the active-standby pair is the default active ServerIron)	10-8
Globally enable FWLB	
Globally enable FWLB	10-8

Configuring a Switch Trunk Group for the Firewall Ports

When you configure FWLB for Layer 2 firewalls, you must configure all the ServerIron ports that are connected to firewalls together as a switch trunk group.

NOTE: To place a trunk group configuration into effect, you must save the configuration to the startup-config file, then reload the software. You can perform these steps as soon as you configure the trunk group or later, after you complete the other firewall configuration steps. In either case, use the **write memory** command to save the configuration to the startup-config file, then enter the **reload** command at the Privileged EXEC level of the CLI to reload the software.

USING THE CLI

To configure a trunk group, enter a command such as the following at the global CLI level:

```
ServerIron(config)# trunk switch ethernet 1 to 2
```

Syntax: [no] trunk switch ethernet <portnum> to <portnum>

You can specify up to four ports. For complete trunk group configuration rules and guidelines, see the "Trunks" section of the *ServerIron TrafficWorks Switching and Routing Guide*.

Specifying the Partner Port

If you are configuring the ServerIron for IronClad FWLB, you need to specify the port number of the dedicated link between the ServerIron and its partner.

USING THE CLI

To specify the port, enter a command such as the following at the global CLI level:

```
ServerIron(config)# server fw-port 13
```

Syntax: [no] server fw-port <portnum>

The command shown above configures port 13 as the dedicated link to the other ServerIron in the active-standby pair. In the example in Figure 10.1 on page 10-2, each of the ServerIrons is configured so that port 13 is the dedicated link to the other ServerIron in the active-standby pair. Thus, the command shown above is entered on each of the ServerIrons. You must specify the partner port on each ServerIron, but using the same port number on each ServerIron is not required.

If the link between the two ServerIrons is a trunk group (recommended for added redundancy), specify the port number of the primary port. The primary port is the first port in the trunk group.

Specifying the Router Ports

IronClad FWLB configurations require paths to the routers as part of the active-standby configuration for the ServerIrons. You need to identify the ports on the ServerIron that are attached to the router(s).

USING THE CLI

To identify port 9 on a ServerIron as a router port, enter the following command:

```
ServerIron(config)# server router-port 9
```

Syntax: [no] server router-ports <portnum>

The command in this example configures port 9 as the router port. In the example shown in Figure 10.1 on page 10-2, Active ServerIron-A and Active ServerIron-B are connected to their router by port 9. The same port number is used for simplicity in this example but you do not need to use the same port number on both ServerIrons.

NOTE: To define multiple router ports on a switch, enter the port numbers, separated by blanks. You can enter up to eight router ports in a single command line. To enter more than eight ports, enter the **server router-port** command again with the additional ports.

Defining the Firewalls and Adding Them to the Firewall Group

When FWLB is enabled, all the ServerIron ports are in firewall group 2 by default. However, you need to add an entry for each firewall. To add an entry for a firewall, specify the firewall name and IP address. You can specify a name up to 32 characters long. After you define the firewalls, add them to the firewall group.

To define the firewalls shown in Figure 10.1 on page 10-2, use the following method.

NOTE: In the case of Layer 2 firewalls, the first part of the firewall name must be the ServerIron port number that is attached to the firewall. In the example in Figure 9.1, the port numbers are 01 and 02. You can use firewall names such as 01fw1 and 02fw2, but not fw1 and fw2.

NOTE: In case of ServerIron chassis devices, both slot and port numbers must be included in the firewall name. For example, if Layer 2 firewalls are attached to a ServerIron on interfaces 3/1 and 3/2, the firewall names can be 03/01fw1 and 03/02fw2.

NOTE: For slot numbers 1 through 8 and port numbers 1 through 9, you must use 0 in the number. For example, 03/01fw1 is a valid name, but 3/1fw1 is not.

USING THE CLI

To define the firewalls using the CLI, enter the following commands:

Commands for Active ServerIron A (External Active)

```
SI-ActiveA(config)# server fw-name 01fw1 1.1.1.100
SI-ActiveA(config-rs-01fw1)# exit
SI-ActiveA(config)# server fw-name 02fw2 1.1.1.101
SI-ActiveA(config-rs-02fw2)# exit
SI-ActiveA(config)# server fw-group 2
SI-ActiveA(config-tc-2)# fw-name 01fw1
SI-ActiveA(config-tc-2)# fw-name 02fw2
```


Commands for Standby ServerIron A (External Standby)

```
SI-StandbyA(config)# server fw-name 01fw1 1.1.1.100
SI-StandbyA(config-rs-01fw1)# exit
SI-StandbyA(config)# server fw-name 02fw2 1.1.1.101
SI-StandbyA(config-rs-02fw2)# exit
SI-StandbyA(config)# server fw-group 2
SI-StandbyA(config-tc-2)# fw-name 01fw1
SI-StandbyA(config-tc-2)# fw-name 02fw2
```

Commands for Active ServerIron B (Internal Active)

```
SI-ActiveB(config)# server fw-name 01fw1 1.1.1.100
SI-ActiveB(config-rs-01fw1)# exit
SI-ActiveB(config)# server fw-name 02fw2 1.1.1.101
SI-ActiveB(config-rs-02fw2)# exit
SI-ActiveB(config)# server fw-group 2
SI-ActiveB(config-tc-2)# fw-name 01fw1
SI-ActiveB(config-tc-2)# fw-name 02fw2
```

Commands for Standby ServerIron B (Internal Standby)

```
SI-StandbyB(config)# server fw-name 02fw1 1.1.1.100
SI-StandbyB(config-rs-01fw1)# exit
SI-StandbyB(config)# server fw-name 02fw2 1.1.1.101
SI-StandbyB(config-rs-02fw2)# exit
SI-StandbyB(config)# server fw-group 2
SI-StandbyB(config-tc-2)# fw-name 02fw1
SI-StandbyB(config-tc-2)# fw-name 02fw2
```

Command Syntax

Syntax: [no] server fw-name <string> <ip-addr>

NOTE: When you add a firewall name, the CLI level changes to the Firewall level. This level is used when you are configuring stateful FWLB.

Syntax: server fw-group 2

This command changes the CLI to firewall group configuration level. The firewall group number is 2. Only one firewall group is supported.

Syntax: [no] fw-name <string>

Adds a configured firewall to the firewall group.

Enabling the L2-fwall Option

For a static route configuration such as the one in the example in 8-1, you need to enable the L2-fwall option on each ServerIron.

USING THE CLI

To enable the L2-fwall option on a ServerIron, enter the following commands:

```
ServerIron(config)# server fw-group 2
ServerIron(config-tc-2)# l2-fwall
```

Syntax: [no] l2-fwall

Configuring Paths and Adding Static MAC Entries for Layer 2

Firewalls

A path is configuration information the ServerIron uses to ensure that a given source and destination IP pair is always authenticated by the same Layer 2 firewall.

Each path consists of the following parameters:

- The path ID – A number that identifies the path. In basic FWLB configurations, the paths go from one ServerIron to the other through the firewalls. The paths go from one ServerIron to the ServerIrons in the other active-standby pair other through the firewalls. A path also goes to the router.
- The ServerIron port – The number of the port that connects the ServerIron to the firewall.
- The other ServerIron's or Layer 2 switch's IP address – The management address of the ServerIron or Layer 2 switch on the other side of the firewall. The ServerIron on the private network side and the other ServerIron or Layer 2 switch are the end points of the data path through the firewall.
- The next hop IP address – Since these are Layer 2 firewalls, the next hop is not an IP interface on the firewall itself, but is instead the same as the destination IP address of the path.

For each type of firewall (Layer 3 synchronous and asynchronous, with or without NAT, or Layer 2), you must configure paths between the ServerIrons through the firewalls.

In addition to configuring the paths, you need to create a static MAC entry for each firewall MAC address.

NOTE: FWLB paths must be fully meshed. When you configure a FWLB path on a ServerIron, make sure you also configure a reciprocal path on the ServerIron attached to the other end of the firewalls.

NOTE: The static MAC entries are required. You must add a static MAC entry for each firewall.

To configure a path and add static MAC entries, use one of the following methods.

USING THE CLI

To configure the paths and static MAC entries for the configuration shown in Figure 10.1 on page 10-2, enter the following commands. Enter the first group of commands on ServerIron A. Enter the second group of commands on ServerIron B.

Commands for Active ServerIron A (External Active)

```
SI-ActiveA(config)# server fw-group 2
SI-ActiveA(config-tc-2)# fwall-info 1 1 1.1.1.30 1.1.1.30
SI-ActiveA(config-tc-2)# fwall-info 2 2 1.1.1.30 1.1.1.30
SI-ActiveA(config-tc-2)# fwall-info 3 1 1.1.1.40 1.1.1.40
SI-ActiveA(config-tc-2)# fwall-info 4 2 1.1.1.40 1.1.1.40
SI-ActiveA(config-tc-2)# fwall-info 5 9 1.1.1.1 1.1.1.1
SI-ActiveA(config-tc-2)# exit
SI-ActiveA(config)# static-mac-address 00e0.5200.3489 ethernet 1 high-priority
router-type
SI-ActiveA(config)# static-mac-address 00e0.5202.e282 ethernet 2 high-priority
router-type
```

Commands for Standby ServerIron A (External Standby)

```
SI-StandbyA(config)# server fw-group 2
SI-StandbyA(config-tc-2)# fwall-info 1 1 1.1.1.30 1.1.1.30
SI-StandbyA(config-tc-2)# fwall-info 2 2 1.1.1.30 1.1.1.30
SI-StandbyA(config-tc-2)# fwall-info 3 1 1.1.1.30 1.1.1.40
SI-StandbyA(config-tc-2)# fwall-info 4 2 1.1.1.30 1.1.1.40
SI-StandbyA(config-tc-2)# fwall-info 5 17 1.1.1.1 1.1.1.1
SI-StandbyA(config-tc-2)# exit
SI-StandbyA(config)# static-mac-address 00e0.5200.3489 ethernet 1 high-priority
router-type
```



```
SI-StandbyA(config)# static-mac-address 00e0.5202.e282 ethernet 2 high-priority
router-type
```

Commands for Active ServerIron B (Internal Active)

```
SI-ActiveB(config)# server fw-group 2
SI-ActiveB(config-tc-2)# fwall-info 1 1 1.1.1.10 1.1.1.10
SI-ActiveB(config-tc-2)# fwall-info 2 2 1.1.1.20 1.1.1.20
SI-ActiveB(config-tc-2)# fwall-info 3 1 1.1.1.10 1.1.1.10
SI-ActiveB(config-tc-2)# fwall-info 4 2 1.1.1.20 1.1.1.20
SI-ActiveB(config-tc-2)# fwall-info 5 9 1.1.1.2 1.1.1.2
SI-ActiveB(config-tc-2)# exit
SI-ActiveB(config)# static-mac-address 00e0.5200.3489 ethernet 1 high-priority
router-type
SI-ActiveB(config)# static-mac-address 00e0.5202.e282 ethernet 2 high-priority
router-type
```

Commands for Standby ServerIron B (Internal Standby)

```
SI-StandbyB(config)# server fw-group 2
SI-StandbyB(config-tc-2)# fwall-info 1 1 1.1.1.10 1.1.1.10
SI-StandbyB(config-tc-2)# fwall-info 2 2 1.1.1.20 1.1.1.20
SI-StandbyB(config-tc-2)# fwall-info 3 1 1.1.1.10 1.1.1.10
SI-StandbyB(config-tc-2)# fwall-info 4 2 1.1.1.20 1.1.1.20
SI-StandbyB(config-tc-2)# fwall-info 5 17 1.1.1.2 1.1.1.2
SI-StandbyB(config-tc-2)# exit
SI-StandbyB(config)# static-mac-address 00e0.5200.3489 ethernet 1 high-priority
fixed-host
SI-StandbyB(config)# static-mac-address 00e0.5202.e282 ethernet 2 high-priority
fixed-host
```

Command Syntax

Syntax: server fw-group 2

Syntax: [no] fwall-info <path-num> <portnum> <other-ServerIron-ip> <next-hop-ip>

The syntax for adding static MAC entries differs depending on whether you are using a stackable or chassis ServerIron.

Syntax for chassis devices:

Syntax: [no] static-mac-address <mac-addr> ethernet <portnum> [priority <0-7>] [host-type | router-type]

Syntax for stackable devices:

Syntax: static-mac-address <mac-addr> ethernet <portnum> [to <portnum> ethernet <portnum>] [normal-priority | high-priority] [host-type | router-type | fixed-host]

The priority can be 0 – 7 (0 is lowest and 7 is highest) for chassis devices and either normal-priority or high-priority for stackable devices.

The defaults are **host-type** and **0** or **normal-priority**.

NOTE: The static MAC entries are required. You must add a static MAC entry for each firewall interface with the ServerIron.

NOTE: Use the **fixed-host** parameter only for Layer 2 firewall configurations such as the one in this example. The parameter “fixes” the address to the ServerIron port you specify and prevents other ports on the ServerIron from learning it. Use the **router-type** parameter for all other types of FWLB configurations. The **fixed-host** parameter is supported only stackable ServerIrons.

NOTE: If you enter the command at the global CONFIG level, the static MAC entry applies to the default port-based VLAN (VLAN 1). If you enter the command at the configuration level for a specific port-based VLAN, the entry applies to that VLAN and not to the default VLAN.

Configuring the ServerIron Priority

If you are configuring the ServerIron for IronClad FWLB, you need to specify the priority for the firewalls within the firewall group. The priority determines which of the partner ServerIrons that are configured together for IronClad FWLB is the default active ServerIron for the firewalls within the group.

You can specify a priority from 0 – 255. The ServerIron with the higher priority is the default active ServerIron for the firewalls within the firewall group.

USING THE CLI

To configure a ServerIron to be the default active ServerIron for the firewalls in group 2, enter the following commands.

Commands for Active ServerIron A (External Active)

```
SI-ActiveA(config)# server fw-group 2
SI-ActiveA(config-tc-2)# sym-priority 255
```

Commands for Standby ServerIron A (External Standby)

To configure another ServerIron to not be the default active ServerIron for the firewalls in group 2, enter the following commands:

```
SI-StandbyA(config)# server fw-group 2
SI-StandbyA(config-tc-2)# sym-priority 1
```

Commands for Active ServerIron B (Internal Active)

```
SI-ActiveB(config)# server fw-group 2
SI-ActiveB(config-tc-2)# sym-priority 255
```

Commands for Standby ServerIron B (Internal Standby)

```
SI-StandbyB(config)# server fw-group 2
SI-StandbyB(config-tc-2)# sym-priority 1
```

Command Syntax

Syntax: [no] sym-priority <num>

The priority can be from 0 – 255.

NOTE: If you specify 0, the CLI removes the priority. When you save the configuration to the startup-config file, the **sym-priority** command is removed. Use this method to remove the priority. You cannot remove the priority using the **no sym-priority** command.

Enabling FWLB

To enable FWLB, you configure global IP policies. FWLB for TCP and UDP is controlled independently, so you need to configure a separate global IP policy for each type of traffic.

When you enable FWLB for TCP or UDP globally, all ports that are in the firewall group are enabled for FWLB. All ServerIron ports are in firewall group 2 by default. Thus, if you enable FWLB globally, it affects all physical ports unless you remove ports from firewall groups.

NOTE: The user interface allows you to enable FWLB locally instead of globally. However, local policies are not applicable to FWLB. Enable the feature globally.

To enable FWLB globally, use the following method.

USING THE CLI

Enter the following commands at the global CONFIG level to enable FWLB for all TCP and UDP traffic:

```
ServerIron(config)# ip policy 1 fw tcp 0 global
ServerIron(config)# ip policy 2 fw udp 0 global
```

Syntax: [no] ip policy <policy-num> fw tcp | udp 0 global

The <policy-num> value identifies the policy and can be a number from 1 – 64.

Each policy affects TCP or UDP traffic, so you must specify **tcp** or **udp**.

The value 0 following the **tcp | udp** parameter specifies that the policy applies to all ports of the specified type (TCP or UDP). In this command, "0" is equivalent to "any port number". For FWLB, you must specify "0".

NOTE: Generally, the firewall itself performs validation and authentication for the traffic, so allowing the ServerIron to pass all traffic of the specified type (TCP or UDP) to the firewall simplifies configuration.

Configuration Example for FWLB with Layer 2 Firewalls

This section shows the ServerIron CLI commands for implementing the configuration shown in Figure 10.1 on page 10-2. Note that the configuration steps for the ServerIrons are similar to those required for the IronClad configuration shown in 8-1 (Layer 3 firewalls in a static route environment).

Commands on Active ServerIron A (External Active)

```
SI-ActiveA(config)# ip address 1.1.1.10/24
SI-ActiveA(config)# ip default-gateway 1.1.1.1
```

The commands above add a management IP address and default gateway address to the ServerIron. For the configuration in this example, the ServerIron needs to be in only one sub-net, so additional IP addresses are not added. However, the IP address must be in the same sub-net as the ServerIron's interface to the Layer 2 firewalls.

```
SI-ActiveA(config)# no span
```

The **no span** command disables the Spanning Tree Protocol (STP). You must disable STP on all the devices in a Layer 2 FWLB configuration such as the one in this example.

```
SI-ActiveA(config)# vlan 2 by port
SI-ActiveA(config-vlan-2)# untagged ethernet 13 to 14
SI-ActiveA(config-vlan-2)# exit
```

The commands above configure a port-based VLAN (separate Layer 2 broadcast domain) for the dedicated link to the partner ServerIron (the other ServerIron in the active-standby pair). The partner link must be in a separate Layer 2 broadcast domain.

```
SI-ActiveA(config)# trunk switch ethernet 13 to 14
```

The **trunk** command creates a trunk group for the ports that connect this ServerIron to its partner. (These are the ports configured in the separate Layer 2 VLAN above.) Using a trunk group for the link between the active and standby ServerIrons is not required, but using a trunk group adds an additional level of redundancy for enhanced availability. If one of the ports in a trunk group goes down, the link remains intact as long as the other port remains up. Since the trunk group is between two ServerIron switches, make sure you configure a switch trunk group, not a server trunk group.

```
SI-ActiveA(config)# server fw-port 13
```


The **server fw-port** command identifies the port that connects this ServerIron to its partner. If you configure a trunk group for the link between the two partners, specify the first port (the primary port for the group) in the trunk group. On the 8-port, 16-port, and 24-port ServerIrons, you can configure a trunk group with two or four members and the primary ports are the odd-numbered ports.

```
SI-ActiveA(config)# server router-port 9
```

The **server router-port** command identifies the port that connects this ServerIron to the router connected to the other ServerIron in the active-standby pair.

```
SI-ActiveA(config)# server fw-name 01fw1 1.1.1.100
SI-ActiveA(config-rs-01fw1)# exit
SI-ActiveA(config)# server fw-name 02fw2 1.1.1.101
SI-ActiveA(config-rs-02fw2)# exit
```

The **server fw-name** commands add the firewalls to the ServerIron. In the commands above, “fw1” and “fw2” are the firewall names. These names are specific to the ServerIron and do not need to correspond to any name parameters on the firewalls themselves. The IP addresses are the addresses of the firewall interfaces with the ServerIron.

The following command, **l2-fwall**, enables the L2-fwall option. This option blocks the Layer 2 traffic on the standby ServerIrons. If you do not enable this option, Layer 2 traffic can pass through the ServerIrons, causing loops. Layer 3 traffic is automatically blocked on the standby ServerIrons, so you do not need to explicitly block the traffic.

```
SI-ActiveA(config)# server fw-group 2
SI-ActiveA(config-tc-2)# l2-fwall
```

The following commands configure the firewall group. The **server fw-group 2** command changes the focus of the CLI to firewall group 2.

The **sym-priority** command specifies the priority of this ServerIron with respect to the other ServerIron for the firewalls in the firewall group. The priority can be from 0 – 255. The ServerIron with the higher priority is the default active ServerIron for the firewalls within the group.

NOTE: If you specify 0, the CLI removes the priority. When you save the configuration to the startup-config file, the **sym-priority** command is removed. Use this method to remove the priority. You cannot remove the priority using the **no sym-priority** command.

The **fw-name** <firewall-name> command adds the firewalls to the firewall group.

```
SI-ActiveA(config-tc-2)# sym-priority 255
SI-ActiveA(config-tc-2)# fw-name 01fw1
SI-ActiveA(config-tc-2)# fw-name 02fw2
```

The **fwall-info** commands add the paths between this ServerIron and the other ServerIrons through the firewalls. The paths enhance performance by ensuring that a given traffic flow (source and destination IP addresses) always travels through the same firewall. In configurations that use asynchronous firewalls, the paths enhance performance by eliminating excess authentications. In this configuration, each ServerIron has two paths to each of the two firewalls. The fifth path goes to the router.

The paths are required, even if the firewalls are synchronized.

The first parameter with each command is a path ID. The second parameter is the port number of the ServerIron port that connects the ServerIron to the firewall in the path.

The third parameter is the IP address of the ServerIron at the other end of the path or, for paths to routers, the IP address of the router’s interface with the ServerIron. Note that each ServerIron has a path to each of the ServerIrons in the other pair, but does not have a path to its own standby pair.

For Layer 2 firewalls, the fourth parameter is also the IP address of the ServerIron at the other end of the path. Notice that the ServerIron has two paths for each firewall. One of the paths goes to the active ServerIron in the other pair. The other path goes to the standby ServerIron in the pair. In the case of the path to the router, the third and forth parameters always have the same value.

```
SI-ActiveA(config-tc-2)# fwall-info 1 1 1.1.1.30 1.1.1.30
SI-ActiveA(config-tc-2)# fwall-info 2 2 1.1.1.30 1.1.1.30
SI-ActiveA(config-tc-2)# fwall-info 3 1 1.1.1.40 1.1.1.40
```



```
SI-ActiveA(config-tc-2)# fwall-info 4 2 1.1.1.40 1.1.1.40
SI-ActiveA(config-tc-2)# fwall-info 5 9 1.1.1.1 1.1.1.1
SI-ActiveA(config-tc-2)# exit
```

The commands below add static entries to the ServerIron's MAC table for the firewall interfaces. The **high-priority** and **fixed-host** parameters are required.

NOTE: Use the **fixed-host** parameter only for Layer 2 firewall configurations such as the one in this example. The parameter "fixes" the address to the ServerIron port you specify and prevents other ports on the ServerIron from learning it. Use the **router-type** parameter for all other types of FWLB configurations. The **fixed-host** parameter is supported only stackable ServerIrons.

```
SI-ActiveA(config)# vlan 1
SI-ActiveA(config-vlan-1)# static-mac-address 00e0.5200.3489 ethernet 1 high-
priority fixed-host
SI-ActiveA(config-vlan-1)# static-mac-address 00e0.5202.e282 ethernet 2 high-
priority fixed-host
SI-ActiveA(config-vlan-1)# exit
```

NOTE: If you enter the command at the global CONFIG level, the static MAC entry applies to the default port-based VLAN (VLAN 1). If you enter the command at the configuration level for a specific port-based VLAN, the entry applies to that VLAN and not to the default VLAN.

The commands below globally enable firewall balancing. The "0" parameter is required and enables the ServerIron to provide FWLB for all packets of the specified type (TCP or UDP). The **write memory** command saves the configuration changes made by all these commands to the ServerIron's startup-config file.

```
SI-ActiveA(config)# ip policy 1 fw tcp 0 global
SI-ActiveA(config)# ip policy 2 fw udp 0 global
SI-ActiveA(config)# write memory
```

Commands on Standby ServerIron **A (External Standby)**

```
SI-StandbyA(config)# ip address 1.1.1.20/24
SI-StandbyA(config)# ip default-gateway 1.1.1.1
SI-StandbyA(config)# no span
SI-StandbyA(config)# vlan 2 by port
SI-StandbyA(config-vlan-2)# untagged ethernet 13 to 14
SI-StandbyA(config-vlan-2)# exit
SI-StandbyA(config)# trunk switch ethernet 13 to 14
SI-StandbyA(config)# server fw-port 13
SI-StandbyA(config)# server router-port 17
SI-StandbyA(config)# server fw-group 2
SI-StandbyA(config-tc-2)# l2-fwall
SI-StandbyA(config-tc-2)# exit
SI-StandbyA(config)# server fw-name 01fw1 1.1.1.100
SI-StandbyA(config-rs-01fw1)# exit
SI-StandbyA(config)# server fw-name 02fw2 1.1.1.101
SI-StandbyA(config-rs-02fw2)# exit
SI-StandbyA(config)# server fw-group 2
SI-StandbyA(config-tc-2)# sym-priority 1
SI-StandbyA(config-tc-2)# fw-name 01fw1
SI-StandbyA(config-tc-2)# fw-name 02fw2
SI-StandbyA(config-tc-2)# fwall-info 1 1 1.1.1.30 1.1.1.30
SI-StandbyA(config-tc-2)# fwall-info 2 2 1.1.1.30 1.1.1.30
SI-StandbyA(config-tc-2)# fwall-info 3 1 1.1.1.40 1.1.1.40
SI-StandbyA(config-tc-2)# fwall-info 4 2 1.1.1.40 1.1.1.40
SI-StandbyA(config-tc-2)# fwall-info 5 17 1.1.1.1 1.1.1.1
SI-StandbyA(config-tc-2)# exit
SI-StandbyA(config)# vlan 1
```



```
SI-StandbyA(config-vlan-1)# static-mac-address 00e0.5200.3489 ethernet 1 high-
priority fixed-host
SI-StandbyA(config-vlan-1)# static-mac-address 00e0.5202.e282 ethernet 2 high-
priority fixed-host
SI-StandbyA(config-vlan-1)# exit
SI-StandbyA(config)# ip policy 1 fw tcp 0 global
SI-StandbyA(config)# ip policy 2 fw udp 0 global
SI-StandbyA(config)# write memory
```

Commands on Active ServerIron **B (Internal Active)**

```
SI-ActiveB(config)# ip address 1.1.1.40/24
SI-ActiveB(config)# ip default-gateway 1.1.1.2
SI-ActiveB(config)# no span
SI-ActiveB(config)# vlan 2 by port
SI-ActiveB(config-vlan-2)# untagged ethernet 13 to 14
SI-ActiveB(config-vlan-2)# exit
SI-ActiveB(config)# server router-port 9
SI-ActiveB(config)# trunk switch ethernet 13 to 14
SI-ActiveB(config)# server fw-port 13
SI-ActiveB(config)# server fw-group 2
SI-ActiveB(config-tc-2)# l2-fwall
SI-ActiveB(config-tc-2)# exit
SI-ActiveB(config)# server fw-name 01fw1 1.1.1.100
SI-ActiveB(config-rs-01fw1)# exit
SI-ActiveB(config)# server fw-name 02fw2 1.1.1.101
SI-ActiveB(config-rs-02fw2)# exit
SI-ActiveB(config)# server fw-group 2
SI-ActiveB(config-tc-2)# sym-priority 255
SI-ActiveB(config-tc-2)# fw-name 01fw1
SI-ActiveB(config-tc-2)# fw-name 02fw2
SI-ActiveB(config-tc-2)# fwall-info 1 1 1.1.1.10 1.1.1.10
SI-ActiveB(config-tc-2)# fwall-info 2 2 1.1.1.20 1.1.1.20
SI-ActiveB(config-tc-2)# fwall-info 3 1 1.1.1.10 1.1.1.10
SI-ActiveB(config-tc-2)# fwall-info 4 2 1.1.1.20 1.1.1.20
SI-ActiveB(config-tc-2)# fwall-info 5 9 1.1.1.2 1.1.1.2
SI-ActiveB(config-tc-2)# exit
SI-ActiveB(config)# vlan 1
SI-ActiveB(config-vlan-1)# static-mac-address 00e0.5200.3490 ethernet 1 high-
priority fixed-host
SI-ActiveB(config-vlan-1)# static-mac-address 00e0.5202.e283 ethernet 2 high-
priority fixed-host
SI-ActiveB(config-vlan-1)# exit
SI-ActiveB(config)# ip policy 1 fw tcp 0 global
SI-ActiveB(config)# ip policy 2 fw udp 0 global
SI-ActiveB(config)# write memory
```

Commands on Standby ServerIron **B (Internal Standby)**

```
SI-StandbyB(config)# ip address 1.1.1.30 255.255.255.0
SI-StandbyB(config)# ip default-gateway 1.1.1.2
SI-StandbyB(config)# no span
SI-StandbyB(config)# vlan 2 by port
SI-StandbyB(config-vlan-2)# untagged ethernet 13 to 14
SI-StandbyB(config-vlan-2)# exit
SI-StandbyB(config)# trunk switch ethernet 13 to 14
SI-StandbyB(config)# server fw-port 13
SI-StandbyB(config)# server router-port 17
```



```
SI-StandbyB(config)# server fw-group 2
SI-StandbyB(config-tc-2)# l2-fwall
SI-StandbyB(config-tc-2)# exit
SI-StandbyB(config)# server fw-name 01fw1 1.1.1.100
SI-StandbyB(config-rs-01fw1)# exit
SI-StandbyB(config)# server fw-name 02fw2 1.1.1.101
SI-StandbyB(config-rs-02fw2)# exit
SI-StandbyB(config)# server fw-group 2
SI-StandbyB(config-tc-2)# sym-priority 1
SI-StandbyB(config-tc-2)# fw-name 01fw1
SI-StandbyB(config-tc-2)# fw-name 02fw2
SI-StandbyB(config-tc-2)# fwall-info 1 1 1.1.1.10 1.1.1.10
SI-StandbyB(config-tc-2)# fwall-info 2 2 1.1.1.20 1.1.1.20
SI-StandbyB(config-tc-2)# fwall-info 3 1 1.1.1.10 1.1.1.10
SI-StandbyB(config-tc-2)# fwall-info 4 2 1.1.1.20 1.1.1.20
SI-StandbyB(config-tc-2)# fwall-info 5 17 1.1.1.2 1.1.1.2
SI-StandbyB(config-tc-2)# exit
SI-StandbyB(config)# vlan 1
SI-StandbyB(config-vlan-1)# static-mac-address 00e0.5200.3490 ethernet 1 high-
priority fixed-host
SI-StandbyB(config-vlan-1)# static-mac-address 00e0.5202.e283 ethernet 2 high-
priority fixed-host
SI-StandbyB(config-vlan-1)# exit
SI-StandbyB(config)# ip policy 1 fw tcp 0 global
SI-StandbyB(config)# ip policy 2 fw udp 0 global
SI-StandbyB(config)# write memory
```


Appendix A
Additional Firewall Configurations

This appendix describes how to configure the following additional firewall configurations:

Configuring FWLB for Firewalls with Active-Standby NICs

Some firewalls provide reliability through link redundancy. For example, some firewalls can have two NICs on each sub-net. One of the NICs is active. The other NIC is a standby interface and is used only if the active NIC becomes unavailable. Both NICs have the same IP address. You can use this type of firewall in IronClad configurations that use the always-active feature.

NOTE: The always-active feature enables you to simplify FWLB configuration by eliminating extra layers of Layer 2 switches. See.

To configure a ServerIron to load balance traffic for firewalls that use dual NICs for link redundancy, specify a wildcard value (255) instead of a specific ServerIron port number when you configure the paths to the firewall. When you add a firewall path, the ServerIron sends an ARP request to obtain the MAC address of the next-hop IP address for the path, which in most configurations is the firewall NIC. If the ServerIron port number for the path is a wildcard (255), the ServerIron also learns the port for the path, which is the port on which the ServerIron receives the ARP reply from the NIC.

Figure 10.2 shows an example of an always-active configuration.

This configuration and the commands for implementing it are almost the same as for the configuration in..... The only differences are as follows:

- Each firewall is connected to both ServerIrons on each side of the network. For example, firewall FW1 is connected to both ServerIron SI-Ext-A and ServerIron SI-Ext-B. Each link has a unique MAC address but they use the same IP address. Only one of the links is active at a time. The other link is a standby.
- The firewall paths on each ServerIron use a wildcard value (255) instead of a specific ServerIron port number.

Figure 10.2 FWLB Configuration Using Always-Active with Active-Standby Firewall Interfaces

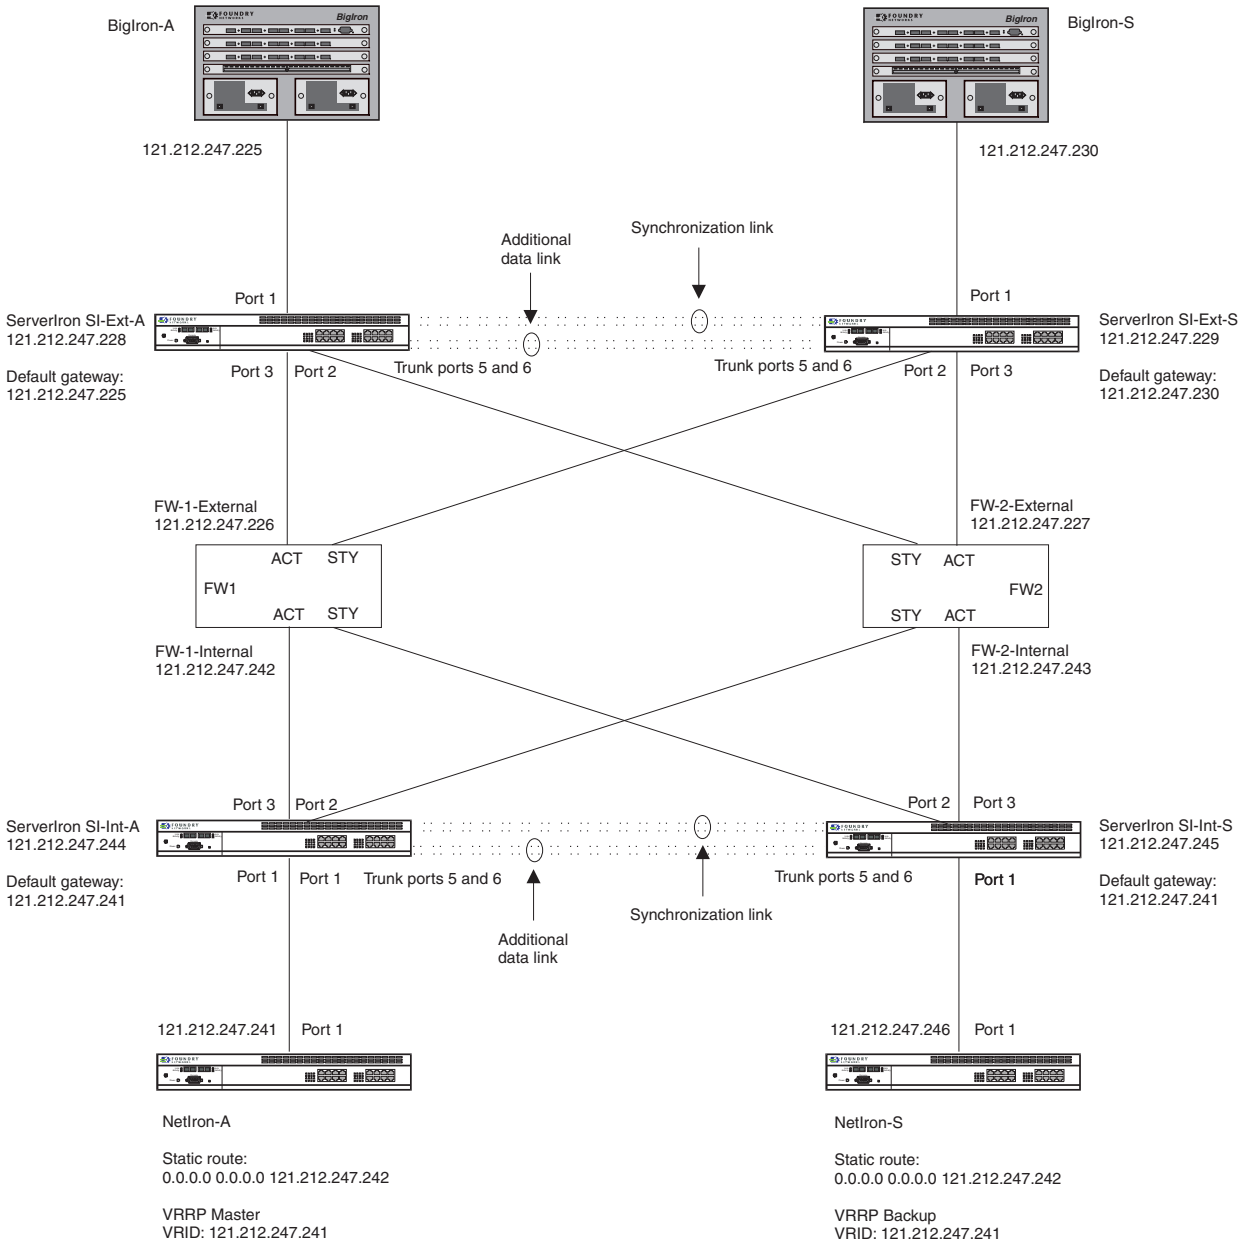

In this example, the links on each firewall are marked to indicate whether they are in the active (ACT) or standby (STY) state. The ServerIron sends traffic to the active firewall interface but not to the standby interface. For example, ServerIron SI-Ext-A sends traffic to firewall FW1 through port 3 because the firewall's link with the ServerIron is on port 3. However, if the link becomes unavailable and the firewall fails over to the other link, ServerIron SI-Ext-A can no longer reach the firewall through port 3. ServerIron SI-Ext-A must use the additional data link configured on ports 5 and 6 (a trunk group in this configuration) to reach the firewall, by sending the traffic through ServerIron SI-Ext-B. (The always-active feature enables the ServerIrons in the active-standby pair to use each other as data paths in instances such as this.)

The ServerIron has only one path to each firewall, but the path uses a wildcard for the ServerIron port number. The ServerIron determines the port to use for reaching the firewall by sending an ARP request for the firewall interface. When the active link on the firewall responds with its MAC address, the ServerIron learns the port on which the response is received and uses that port to reach the firewall.

If the firewall link goes down and the NIC fails over to the other connection, the ServerIron learns the new port for the MAC address. Generally, this occurs when the NIC sends a gratuitous ARP to advertise the new MAC address. The ServerIron learns that the link has failed when the firewall path health check fails. The path health check consists of an IP ping to the next-hop IP address of the path.

Configuring for Active-Standby Firewall Links

To configure firewall paths for firewalls with active-standby NICs, enter commands such as the following. Notice that the first four paths configured for each ServerIron specify 255 as the ServerIron port number (the second parameter in the command). The last path is the path to the router and does use a specific ServerIron port instead of the wildcard (255).

Commands for Active External ServerIron (SI-Ext-A)

```
SI-Ext-A(config)# server fw-group 2
SI-Ext-A(config-tc-2)# fwall-info 1 255 121.212.247.244 121.212.247.226
SI-Ext-A(config-tc-2)# fwall-info 2 255 121.212.247.245 121.212.247.226
SI-Ext-A(config-tc-2)# fwall-info 3 255 121.212.247.244 121.212.247.227
SI-Ext-A(config-tc-2)# fwall-info 4 255 121.212.247.245 121.212.247.227
SI-Ext-A(config-tc-2)# fwall-info 5 1 121.212.247.225 121.212.247.225
```

Commands for Standby External ServerIron (SI-Ext-S)

```
SI-Ext-S(config)# server fw-group 2
SI-Ext-S(config-tc-2)# fwall-info 1 255 121.212.247.244 121.212.247.226
SI-Ext-S(config-tc-2)# fwall-info 2 255 121.212.247.245 121.212.247.226
SI-Ext-S(config-tc-2)# fwall-info 3 255 121.212.247.244 121.212.247.227
SI-Ext-S(config-tc-2)# fwall-info 4 255 121.212.247.245 121.212.247.227
SI-Ext-S(config-tc-2)# fwall-info 5 1 121.212.247.230 121.212.247.230
```

Commands for Active Internal ServerIron (SI-Int-A)

```
SI-Int-A(config)# server fw-group 2
SI-Int-A(config-tc-2)# fwall-info 1 255 121.212.247.228 121.212.247.242
SI-Int-A(config-tc-2)# fwall-info 2 255 121.212.247.229 121.212.247.242
SI-Int-A(config-tc-2)# fwall-info 3 255 121.212.247.228 121.212.247.243
SI-Int-A(config-tc-2)# fwall-info 4 255 121.212.247.229 121.212.247.243
SI-Int-A(config-tc-2)# fwall-info 5 1 121.212.247.241 121.212.247.241
```

Commands for Standby Internal ServerIron (SI-Int-S)

```
SI-Int-S(config)# server fw-group 2
SI-Int-S(config-tc-2)# fwall-info 1 255 121.212.247.228 121.212.247.242
SI-Int-S(config-tc-2)# fwall-info 2 255 121.212.247.229 121.212.247.242
SI-Int-S(config-tc-2)# fwall-info 3 255 121.212.247.228 121.212.247.243
SI-Int-S(config-tc-2)# fwall-info 4 255 121.212.247.229 121.212.247.243
SI-Int-S(config-tc-2)# fwall-info 5 1 121.212.247.246 121.212.247.246
```

Syntax: [no] fwall-info <path-num> <portnum> <other-ServerIron-ip> <next-hop-ip>

Specify 255 as the port number for the paths to dual NIC (active-standby) firewall interfaces. Specify the ServerIron port number for paths to routers.

When the firewalls have active-standby NICs, and dynamic ports are configured on the firewall paths, by default the ServerIron always uses the same interface to reach a firewall, where firewall's ARP entry was initially learnt. It does not update the firewall path to an alternate interface unless the interface physically goes down.

This behavior will cause issues in setups running Firewalls with active-standby NIC's, when the NICs fail over without having the interface go down physically. For example, when a failover of the Firewall NIC occurs, the ARP entry for the firewall's IP is learnt on a new port but the firewall path still shows the old interface causing issues with FWLB.

Configure the following command, to prevent this condition:

```
ServerIron# server fw-allow-dynamic-port-change
```

This command allows the firewall path health checks to be sent to the correct port where the firewall ARP is learnt and update the firewall path accordingly to reflect the new interface where the firewall can now be reached.

NOTE: For the complete CLI example, see.... The example in the Guide does not use the wildcard in the firewall paths and the firewalls do not have active-standby NICS, but the other aspects of the configurations are the same.

Customizing Path Health Checks

This appendix describes the health checks for firewall and router paths and how to change their configuration.

By default, the ServerIron checks the health of each firewall and router path by sending an ICMP ping on the path every 400 milliseconds.

- If the ServerIron receives one or more responses within 1.2 seconds, the ServerIron concludes that the path is healthy.
- Otherwise, the ServerIron reattempts the health check by sending another ping. By default, the ServerIron reattempts an unanswered path health check up to three times before concluding that the path is unhealthy.

You can change the maximum number of retries for the Layer 3 health checks of firewall and router paths. You also can enable Layer 4 path health checks for the firewall paths.

NOTE: This chapter describes how to configure path health checks but not application health checks. To configure a Layer 4 or Layer 7 application health check, use the procedures in the "Configuring Health Checks" section of the "Configuring Port and Health Check Parameters" chapter in the *Foundry ServerIron Installation and Configuration Guide*. To configure a Layer 4 or Layer 7 application health check, use the procedures in the "Health Checks" chapter of the *ServerIron TrafficWorks Server Load Balancing Guide*.

Changing the Maximum Number of Layer 3 Path Health-Check Retries

By default, the ServerIron checks the health of each firewall and router path by sending an ICMP ping on the path every 400 milliseconds.

- If the ServerIron receives one or more responses within 1.2 seconds, the ServerIron concludes that the path is healthy.
- Otherwise, the ServerIron reattempts the health check by sending another ping. By default, the ServerIron reattempts an unanswered path health check up to three times before concluding that the path is unhealthy.

You can change the maximum number of retries to a value from 3 – 31 (ServerIron Chassis devices) or 8 – 31 (all other ServerIron models).

To change the maximum number of FWLB path health check attempts, enter a command such as the following at the firewall level of the CLI:

```
ServerIron(config-tc-2)# fw-health-check icmp 20
```

Syntax: [no] fw-health-check icmp <num>

The <num> parameter specifies the maximum number of retries and can be a number from 3 – 31 (ServerIron Chassis devices) or 8 – 31 (all other ServerIron models). The default is 3 for ServerIron Chassis devices, 8 for all other ServerIron models.

Enabling Layer 4 Path Health Checks for FWLB

By default, the ServerIron performs Layer 3 health checks of firewall paths, but does not perform Layer 4 health checks of the paths. You can configure the ServerIrons in an FWLB configuration to use Layer 4 health checks instead of Layer 3 health checks for firewall paths. When you configure a Layer 4 health check, the Layer 3 (ICMP) health check, which is used by default, is disabled.

NOTE: The Layer 4 health check applies only to firewall paths. The ServerIron always uses a Layer 3 (ICMP) health check to test the path to the router.

When you configure a Layer 4 health check for firewall paths, the ServerIron sends Layer 4 health checks and also responds at Layer 4 to health checks from the ServerIron at the other end of the firewall path.

To configure a Layer 4 health check, specify the protocol (TCP or UDP). Optionally, you also can specify the port.

- UDP – The ServerIron sends and listens for path health check packets on the port you specify. If you do not specify a port, the ServerIron uses port 7777 by default. The port number is used as both the source and destination UDP port number in the health check packets.
- TCP – The ServerIron listens for path health check packets on the port you specify, but sends them using a randomly generated port number. If you do not specify a port, the ServerIron uses port 999 as the destination port by default.

NOTE: You must configure the same Layer 4 health check parameters on all the ServerIrons in the FWLB configuration. Otherwise, the paths will fail the health checks.

To configure a Layer 4 health check for firewall paths, enter a command such as the following at the firewall group configuration level:

```
ServerIron(config-tc-2)# fw-health-check udp
```

The command in this example enables Layer 4 health checks on UDP port 7777. This ServerIron sends firewall path health checks to UDP port 7777 and listens for health checks on UDP port 7777.

Syntax: [no] fw-health-check udp | tcp [<tcp/udp-portnum> <num>]

The <tcp/udp-portnum> parameter specifies the TCP or UDP port and can be a number in one of the following ranges:

- For TCP, from 1 – 65535
- For UDP, from 1 – 1032 or 2033 – 65535

NOTE: Do not use a number from 1033 – 2032 for UDP. Port numbers in this range are not supported for FWLB UDP health checks.

The <num> parameter specifies the maximum number of retries and can be a number from 8 – 31. The default is 3.

Disabling Layer 4 Path Health Checks on Individual Firewalls and Application Ports

To disable the Layer 4 health check for an individual application on an individual firewall, enter a command such as the following at the firewall configuration level of the CLI:

```
ServerIron(config-rs-FW1)# port http no-health-check
```

The command in this example disables Layer 4 health checks for port HTTP on firewall FW1.

Syntax: [no] no-health-check

FWLB Selection Algorithms

This appendix describes selection algorithms for FWLB. This appendix contains the following sections:

- Least Connections
- Least Connections per Application
- Hashing

NOTE: If hash-port is configured, hashing includes both source-port and destination-port.

Hashing Based on Destination TCP or UDP Application Port

The ServerIron uses a hash value based on the source and destination IP addresses in a packet to select a path, and thus a firewall, for the packet. After calculating this hash value for a given source-and-destination pair, the ServerIron always uses the same path and firewall for packets containing that source-and-destination pair.

You can configure the ServerIron to also hash based on TCP or UDP port numbers. This is useful in environments where the same source-and-destination pairs generate a lot of traffic and you want to load balance the traffic across more than one firewall.

For example, if you configure the ServerIron to hash based on TCP ports 69 (TFTP) and 80 (HTTP), the ServerIron hashes packets addressed to one of these ports by calculating a hash value based on the source and destination IP addresses and the TCP port number (69 or 80). Since the TCP port numbers are included in the hash calculations for these packets, the calculations can result in packets for port 80 receiving a different hash value (and thus possibly a different path and firewall) than packets for port 69, even though the source and destination IP addresses are the same.

NOTE: The current release supports stateful FWLB only for TCP/UDP applications that do not require multiple simultaneous connections for the same client to the same firewall. For example, you cannot use stateful FWLB for FTP, because this application requires separate simultaneous control and data connections to the firewall. The CLI allows you to specify FTP or any other port, but you might not receive the desired results if the application uses multiple simultaneous connections to the same firewall.

You can specify a list of ports, a range of ports, or both. The software hashes based on the combined set of ports from the list and the range.

Specifying a List of Application Ports for Use When Hashing

To specify a list TCP/UDP ports to include in the hash calculations, use either of the following methods.

USING THE CLI

To specify a list of TCP/UDP ports for hashing, enter commands such as the following:

```
ServerIron(config)# server fw-group 2
ServerIron(config-tc-2)# hash-ports 69 80
```

Syntax: [no] hash-ports <num> [<num...>]

The <num> parameters specify TCP or UDP port numbers. You can specify up to eight port numbers on the same command line.

Specifying a Range of Application Ports for Use When Hashing

To specify a range of application ports, enter a command such as the following at the firewall group configuration level of the CLI:

```
ServerIron(config-tc-2)# hash-port-range 69 80
```

Syntax: [no] hash-port-range <start-num> <end-num>

The <start-num> parameter specifies the starting port number in the range. Specify the port number at the lower end of the range.

The <end-num> parameter specifies the ending port number in the range. Specify the port number at the higher end of the range.

Overriding the Global Hash Values

By default, the ServerIron uses the hash mask you configure for the firewall group for all hash-based load balancing of firewall traffic. You can override the global hash mask for specific traffic based on source or destination address information.

Here is a CLI example:

```
ServerIron(config)# access-list 100 permit ip any 192.168.1.16 0.0.0.15
ServerIron(config)# access-list 100 permit ip any 192.168.2.0 0.0.0.255
ServerIron(config)# access-list 100 permit ip any 192.168.3.192 0.0.0.63
ServerIron(config)# access-list 100 permit ip any 192.168.4.0 0.0.0.255
ServerIron(config)# access-list 100 permit ip any 192.168.3.160 0.0.0.31
ServerIron(config)# access-list 100 permit ip any 192.168.3.0 0.0.0.127
ServerIron(config)# access-list 100 permit ip any 64.129.1.0 0.0.0.255
ServerIron(config)# server fw-group-2
ServerIron(config-tc-2)# hash-mask 255.255.255.255 0.0.0.0
ServerIron(config-tc-2)# policy-hash-acl 100 255.255.255.255 255.255.255.255
```

In this example, FWLB will use the hash mask 255.255.255.255 0.0.0.0 for all traffic ***except*** the traffic that matches ACL 100.

Syntax: [no] server policy-hash-acl <acl-id> <dst-mask> <src-mask>

The <acl-id> parameter specifies a standard or extended ACL. Configure each entry in the ACL to permit the addresses for which you want to override the global hash mask.

The <dst-mask> parameter species the destination mask.

The <src-mask> parameter species the source mask.

For information about configuring standard and extended ACLs, see the "Access Control List" chapter in the *ServerIron TrafficWorks Security guide*.

Configuring Weighted Load Balancing

You can assign weights to your firewalls, to bias the load balancing in favor of certain firewalls.

Weight

The weight you assign to a firewall determines the percentage of the current connections that are given to that firewall. For example, in a configuration with five firewalls of various weights, the percentage of connections is calculated as follows:

- Weight fwall1 = 7
- Weight fwall2 = 8
- Weight fwall3 = 2
- Weight fwall4 = 2
- Weight fwall5 = 5

Total weight of all firewalls = 24

The result is that fwall1 gets 7/24 of the current number of connections, fwall2 gets 8/24, server3 gets 2/24, and so on. If a new firewall, fwall6, is added with a weight of 10, the new firewall gets 10/34.

If you set the weight so that your fastest firewall gets 50 percent of the connections, it will get 50 percent of the connections at a given time. Because the firewall is faster than others, it can complete more than 50 percent of the total connections overall because it services the connections at a higher rate. Thus, the weight is not a fixed ratio but adjusts to firewall capacity over time.

The default weight for firewalls is 1.

The weight feature is supported only for stateful FWLB. FWLB in software releases 07.2.x and 08.x is always stateful. FWLB in releases 07.1.x and 07.3.x can be stateful or stateless, depending upon your configuration.

Assigning Weights to Firewalls

To assign weights to firewalls, enter commands such as the following:

```
ServerIron(config)# server fw-name fw1
ServerIron(config-rs-fw1)# weight 7
ServerIron(config-rs-fw1)# server fw-name fw2
ServerIron(config-rs-fw2)# weight 8
ServerIron(config-rs-fw2)# server fw-name fw3
ServerIron(config-rs-fw3)# weight 2
ServerIron(config-rs-fw3)# server fw-name fw4
ServerIron(config-rs-fw4)# weight 2
ServerIron(config-rs-fw4)# server fw-name fw5
ServerIron(config-rs-fw5)# weight 5
```

These commands assign weights to five firewalls. The ServerIron will load balance new connections to the firewalls based on their relative weights.

Syntax: [no] weight <least-connections-weight>

The <least-connections-weight> parameter assigns a weight to the firewall. This weight determines the percentage of new connections the firewall receives relative to the other firewalls.

NOTE: The **weight** command has a second parameter, <response-time-weight>. This parameter is valid for real servers in SLB configurations but is not valid for FWLB.

Denying FWLB for Specific Applications

You can deny FWLB for specific applications while still permitting FWLB for other applications. For example, you can deny FWLB for HTTP traffic (TCP port 80) while still providing FWLB for other types of traffic.

This feature is useful when your network is configured to send all traffic for a given application to the same firewall. For example, Figure A.1 shows a network in which the routers are configured to send all HTTP traffic through firewall FW1.

Figure A.1 FWLB Denied for Application Traffic

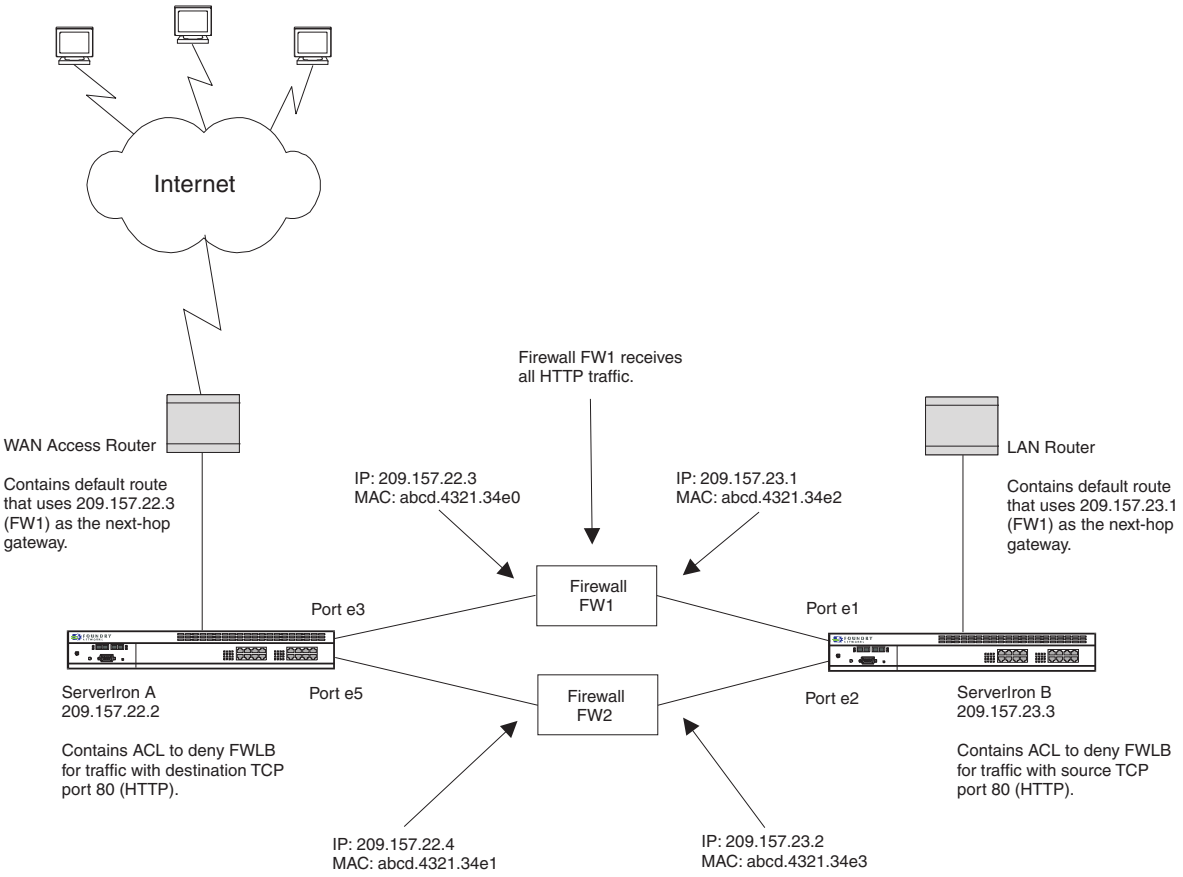

In this example, the network is configured as follows:

- The WAN access router has a default route that identifies IP address 209.157.22.3 on FW1 as the next-hop gateway.
- The LAN router has a default route that identifies IP address 209.157.23.1 (also on FW1) as the next-hop gateway.
- ServerIron A has an extended ACL at the firewall group configuration level that denies FWLB for packets addressed to destination TCP port 80.
- ServerIron B has an extended ACL at the firewall group configuration level that denies FWLB for packets from source TCP port 80.

Notice that the routers use default routes to send traffic to a specific firewall. However, the default routes do not necessarily determine the firewall to which the ServerIron sends the traffic. When the ServerIron performs load balancing for a packet and selects a firewall for the traffic, the ServerIron also changes the destination MAC address of the packet to the MAC address of the firewall selected by the ServerIron. For example, in Figure A.1, if ServerIron A selects firewall FW2 for a packet, the ServerIron changes the destination MAC address of the packet to abcd.4321.34e1,the MAC address of firewall FW2's interface with ServerIron A. As a result, even if the WAN access router addresses a packet to the MAC address for firewall FW1, the ServerIron does not send the packet to firewall FW1 unless the load balancing mechanism selects that firewall. In either case, the ServerIron changes the destination MAC address of the packet.

If you want to ensure that all packets for an application go to a specific firewall (as specified in the default route on the router), you must deny FWLB service for that application. For example, if you have configured firewall FW1 to collect statistics on HTTP traffic and you therefore want to send all the HTTP traffic to firewall FW1, you must

disable FWLB for HTTP traffic. To disable FWLB for an application, configure an extended ACL at the firewall group configuration level.

NOTE: When you configure an ACL at the firewall group configuration level, a deny action does not cause the ServerIron to drop the denied packet. In this type of configuration, a deny action denies FWLB service for the packet, so that the ServerIron leaves the destination MAC address of the packet unchanged.

NOTE: This section focuses on using extended ACLs to deny FWLB based on TCP or UDP port. However, you also can use standard ACLs at the firewall group configuration level to deny FWLB based on IP address.

Configuration Guidelines

- Global IP policies to enable FWLB are still required. You must enable FWLB globally for all TCP traffic and all UDP traffic.
- Configure extended ACLs at the firewall group configuration level to deny FWLB for specific applications.
- Configure a permit ACL to allow all applications. Once you configure an ACL, the default action changes from permit to deny. As a result, if you do not configure the permit ACL for all traffic types, FWLB is denied for all traffic. Make sure the permit ACL for all traffic is the last ACL, after all the deny ACLs.
- Configure the deny ACLs for each direction of traffic for which you want to deny FWLB. In Figure A.1, configure a deny ACL on ServerIron A to deny FWLB for packets addressed to destination TCP port 80 (HTTP). To deny FWLB for the return traffic, configure a deny ACL on ServerIron B to deny packets from source TCP port 80.

Denying FWLB

To deny FWLB for an application, enter commands such as the following. These commands configure the ServerIrons in Figure A.1 to deny FWLB for HTTP traffic, in both directions. On ServerIron A, FWLB is denied for traffic addressed to TCP port 80. On ServerIron B, FWLB is denied for traffic from TCP port 80.

ServerIron A Commands

```
ServerIronA(config)# ip policy 1 fw tcp 0 global
ServerIronA(config)# ip policy 2 fw udp 0 global
ServerIronA(config)# access-list 101 deny tcp any any eq http
ServerIronA(config)# access-list 101 permit tcp any any
ServerIronA(config)# access-list 101 permit udp any any
ServerIronA(config)# server fw-group 2
ServerIronA(config-tc-2)# acl-id 101
```

The first two commands globally enable FWLB for all TCP and UDP applications. These commands are required. The following commands configure three ACL entries. The first entry denies FWLB for packets addressed to TCP port 80 (HTTP). The second ACL permits FWLB for all TCP applications. Packets that do not match the first ACL entry match the second ACL entry and are provided with FWLB. The third ACL permits FWLB for all UDP applications. The last two commands change the CLI level to the firewall group configuration level and apply ACL 101 to the firewall group.

Syntax: [no] access-list <num> deny | permit <ip-protocol> <source-ip> | <hostname> <wildcard> [<operator> <source-tcp/udp-port>] <destination-ip> | <hostname> <wildcard> [<operator> <destination-tcp/udp-port>] [precedence <name> | <num>] [tos <name> | <num>] [log]

Syntax: [no] acl-id <num>

For detailed information about the ACL syntax, see the “Access Control List” chapter in the *ServerIron TrafficWorks Security Guide*.

ServerIron B Commands

```
ServerIronB(config)# ip policy 1 fw tcp 0 global
ServerIronB(config)# ip policy 2 fw udp 0 global
ServerIronB(config)# access-list 101 deny tcp any eq http any
ServerIronB(config)# access-list 101 permit tcp any any
ServerIronB(config)# access-list 101 permit udp any any
ServerIronB(config)# server fw-group 2
ServerIronB(config-tc-2)# acl-id 101
```

These commands are the same as the commands on ServerIron A, except the first ACL entry matches on TCP port 80 (eq http) as the destination TCP port on ServerIron A, but matches as the source TCP port on ServerIron B.

Configuring Failover Tolerance in IronClad Configurations

By default, failover from the active ServerIron to the standby ServerIron in an IronClad configuration occurs if a path link on the active ServerIron becomes unavailable. If all the path links are stable, failover is an uncommon event. However, an unreliable link can cause frequent failover. For example, if a link on a firewall flaps (goes up and down) frequently, the flapping can cause frequent, unnecessary failovers.

You can reduce the frequency of such failovers by specifying a path link tolerance for firewall paths and for router paths. The tolerance specifies the minimum number of such paths that must be good in order for the active ServerIron to remain active. Only if the number of paths is less than the configured minimum and less than the number of available paths on the other ServerIron does failover occur. If the number of paths remains equal on each ServerIron, even if some paths are unavailable on each ServerIron, failover does not occur.

The default failover tolerance for firewall paths is one half the configured firewall paths. The default tolerance for router ports is one half the configured router ports.

To change the minimum number of paths required on a ServerIron, use the following method.

NOTE: The minimum number of required paths must match on each ServerIron in an active-standby pair. For example, if you specify one router path and three firewall paths as the minimum on the active ServerIron, you must configure the same minimums on the standby ServerIron.

USING THE CLI

To specify the minimum number of paths required on a ServerIron, enter commands such as the following:

```
ServerIron(config)# server fw-group 2
ServerIron(config-tc-2)# prefer-router-cnt 1
ServerIron(config-tc-2)# prefer-cnt 3
```

This example specifies that a minimum of one router path and three firewall paths must be available for the ServerIron to remain active. Thus, if the ServerIron has four firewall paths, one path can be unavailable and the ServerIron will remain the active ServerIron.

Syntax: [no] prefer-router-cnt <num>

Syntax: [no] prefer-cnt <num>

For each command, the <num> parameter specifies the minimum number of paths required.

